Pseudofinite Model Theory

Anand Pillay, and friends

April 13, 2015

These notes are based on a course given by Anand Pillay in the Autumn of 2014 at the University of Notre Dame. They were transcribed from the lectures by Matteo Bianchetti, Gregory Cousins, Michael Haskel, Paul McEldowney, Nathan Pierson, Somayeh Vojdani, and Rose Weisshaar. The current set of notes were compiled and edited by Gregory Cousins and Michael Haskel. The notes are in a rough state, and therefore likely contain many inconsistencies, mistakes, typos, questionable arguments, and missing information. Please feel free to send corrections and suggestions to Greg Cousins at gcousins@nd.edu.

Contents

Chapter I

Definitions and Basic Examples

We assume the reader is familiar with basic definitions, notations, and results from model theory.

Definition I.1. The basic definitions of pseudofiniteness are as follows:

- An L-structure M is *pseudofinite* if for all L-sentences σ, $M \models \sigma$ implies that there is a finite M_0 such that $M_0 \models \sigma$. M is *strictly pseudofinite* if M is pseudofinite and not finite.
- A consistent L-theory T is *weakly pseudofinite* if whenever $T \models \sigma$, then σ is true in some finite structure (not necessarily a model of T).
- T is *strongly pseudofinite* if whenever σ is consistent with T, there exists a finite model M_0 such that $M_0 \models \sigma$.

For example, the empty theory is weakly pseudofinite but not strongly pseudofinite.

Definition I.2. Fixing a language L, T_f is the common theory of all finite L-structures. That is, $\sigma \in T_f$ if and only if σ is true of every finite L-structure.

There are other, equivalent definitions of pseudofiniteness, as the following result describes:

Proposition I.3. *Fix a language L and an L-structure M. Then the following are equivalent:*

1. *An L-structure M is pseudofinite.*
2. $M \models T_f$
3. *M is elementarily equivalent to an ultraproduct of finite L-structures.*

Proof. $(1 \Leftrightarrow 2)$ M is fails to be pseudofinite iff there is σ such that $M \models \sigma$, with σ not true in any finite structure. That is, $M \models \sigma$ with $\neg\sigma \in T_f$. That is, $M \not\models T_f$.

$(3 \Rightarrow 2)$ It is sufficient to show that any ultraproduct of finite structures models T_f. Łos' Theorem provides this result, since each of the finite structures in such an ultraproduct model T_f.

$(2 \Rightarrow 3)$ Let T be a complete theory extending T_f. It is sufficient to build a model of T as an ultraproduct of finite structures. The following argument is essentially the ultraproduct proof of the compactness theorem, modified to use only finite structures. This modification is possible exactly because T extends T_f.

Consider the boolean algebra $\mathcal{B}$ of L-sentences under equvalence modulo T_f. Since T extends T_f, we can interpret T as an ultrafilter of $\mathcal{B}$. Additionally, let $\mathcal{S}$ be the set of finite L-structures, and $\mathcal{P}(\mathcal{S})$ be the boolean algebra of subsets of $\mathcal{S}$.

Define $g : \mathcal{B} \to \mathcal{P}(\mathcal{S})$ to be the map sending a sentence to the set of structures in $\mathcal{S}$ modeling it. The map g is well defined and injective because two formulas are equivalent module T_f iff they are true on the same set of finite structures. The map g is trivially a boolean algebra homomorphism.

By general facts about boolean algebra, since T is a proper filter of $\mathcal{B}$ and g is an injective homomorphism, $g(T)$ is a proper filter of $\mathcal{P}(\mathcal{S})$. Take $\mathcal{U}$ an ultrafilter extending $g(T)$, and let M be the ultraproduct of $\mathcal{S}$ along $\mathcal{U}$. We obtain $M \models T$, as desired, by Łos' Theorem.

□

Exercise I.4. Recall that an L-theory T is weakly pseudofinite if every sentence σ in T is true in some finite model. Prove the following:

1. Show that T is weakly pseudofinite iff $T \cup T_f$ has some pseudofinite model.
2. Show T is strongly pseudofinite iff $T \models T_f$, which is true if and only if every model of T is pseudofinite.
3. Show that, if T is complete, then T is weakly pseudofinite iff T is strongly pseudofinite.

In light of the previous exercise, whenever T is a complete theory, we say T is pseudofinite to mean either weakly or strongly pseudofinite.

Example I.5. 1. For a language L, the common theory of all L-structures is weakly pseudofinite.

2. Let T_0 be any L-theory. Let $T_{0,f}$ be the common first-order theory of all finite models of T_0. Then $T_{0,f}$ is strongly pseudofinite. Note that $T_{0,f}$ may not be the weakest strongly pseudofinite theory extending T_0.

3. Fix a finite field F, and let T be the theory of F-vector spaces in the language of additive groups. Then T is strongly pseudofinite.

   Proof. Let T' be the theory of infinite models of T. It is sufficient to show that T' is strongly pseudofinite. It is clear that T' is weakly pseudofinite, since every finite conjunction of the natural axiomatization of T' has a finite model. Furthermore, by an easy application of Vaught's test, T' is complete, and so consequently T' is strongly pseudofinite. $\square$

4. The theory T of $\mathbb{Q}$-vector spaces in the language of groups with a function symbol for scalar multiplication by each rational is complete and pseudofinite.

   Proof. Axiomatize T with the axioms for Abelian groups, together with a sentence $\sigma_{\frac{r}{s}}$ for each $\frac{r}{s} \in \mathbb{Q}$, where $\sigma_{\frac{r}{s}}$ asserts

   $$(\forall x) \sum_{i=1}^{s} \left(\frac{r}{s}\right) \cdot x = \sum_{i=r}^{r} x.$$

   As in the previous example, this is a complete axiomatization for T by an easy application of Vaught's test (this time for some uncountable cardinality). Likewise, T is weakly pseudofinite because we can satisfy any finite conjunction of axioms by a finite model. $\square$

5. The theory of dense linear orders is not pseudofinite.

6. The theory ACF of algebraically closed fields is not weakly pseudofinite.

   Proof. Consider that, in an algebraically closed field of characteristic other than 2, the map sending x to x^2 is surjective but not injective. If the characteristic is 2, then the map sending x to x^3 has this property. Let σ be the sentence asserting "I am a field, and either the square map or the cube map is surjective but not injective." We have seen that $ACF \models \sigma$. However, σ has no finite model, since a function from a finite set to itself is injective iff it is surjective. $\square$

Next, we describe the key construction used in the study of pseudofinite theories. The general idea is that, in finite theories, it is possible to measure the size of sets as natural numbers. These counting measures transfer to arbitrary pseudofinite models via the following constructions.

Definition I.6. Let $\mathcal{R}$ be the real numbers, together with both their usual algebraic structure and a predicate for the natural numbers. Take $\mathcal{R}^*$ to be a sufficiently saturated elementary extension of $\mathcal{R}$. Let $\mathcal{N}$ and $\mathcal{N}^*$ denote the (definable) set of natural numbers in $\mathcal{R}$ and $\mathcal{R}^*$, respectively. We call elements of $\mathcal{R}^* \setminus \mathcal{R}$ nonstandard numbers.

Note that, while $\mathcal{N}^*$ is not necessarily well-ordered, every definable subset (with parameters) has a least element. More generally, any fact about the natural numbers (or real numbers) that we can state in a first order way holds of $\mathcal{N}^*$ and $\mathcal{R}^*$ as well. We will call this ad hoc principle the *transfer* principle and use it frequently.

Definition I.7. Let $\mathrm{Fin}(\mathcal{R}^*)$ be the set of *finite* nonstandard reals, that is, those $x \in \mathcal{R}^*$ such that $|x| < n$ for some *standard* natural number n. Note that $\mathrm{Fin}(\mathcal{R}^*)$ is not definable, or even type definable, even with parameters.

For finite reals, we have the following fact:

Fact I.8. *There is a standard part map* $\mathrm{st} : \mathrm{Fin}(\mathcal{R}^*) \to \mathcal{R}$ *where* $\mathrm{st}(a)$ *is the unique standard real such that* $|\,\mathrm{st}(a) - a| < \frac{1}{n}$ *for all standard* $n \in \mathcal{N}$.

The standard part map behaves nicely with respect to orderings, algebra, etc. We leave the details of its behavior as an exercise for the reader.

We now describe the details of how to measure sizes of sets in pseudofinite structures. What follows are two equivalent constructions. The first uses ultraproducts, and the second uses compactness directly.

Construction I.9. Let L be a language with home sort M. Let T be a complete pseudofinite L-theory. Let L' be L expanded with the following:

- an additional sort $\mathcal{R}^*$ for real numbers, together with their algebraic structure and predicate for the natural numbers;
- for each $n \in \omega$, an additional "powerset" sort $\mathcal{P}^*(M^n)$, together with appropriate "$\in$" symbols; and
- for each $n \in \omega$, a "cardinality" function symbol from $\mathcal{P}^*(M^n)$ to $\mathcal{R}^*$.

Since T is pseudofinite, take an L-structure $M \models T$, with

$$M = \prod_{\mathcal{U}} M_i,$$

where the M_i are finite L-structures. It is clear that each of the finite M_i expands canonically to an L'-structure M_i': the new sorts are the standard reals and standard powersets, and the cardinality symbols measure standard finite cardinalities.

Define

$$M' = \prod_{\mathcal{U}} M_i'$$

an expansion of M to an L'-structure.

Note that the theory of M' does not depend only on T. In general, the theory of M' might depend on which collecion of M_i and which ultrafilter we choose. Let T' be the (possibly incomplete) L'-theory consisting of the L'-sentences true in every possible M'.

Note that, while T' in Construction I.9 is possibly incomplete, it contains enough information (by Łos' theorem) to argue via the transfer principle about how counting and subsets work in finite structures. The following, equivalent construction makes this fact explicit.

Construction I.10. L, and L' be as in Construction I.9. As in Construction I.9, observe that every finite L-structure expands canonically to an L'-structure. Recalling that T_f is the common L-theory of finite L-structures, define T'_f to be the common L'-theory of all such canonically expanded finite L-structures. It is immediate that this expansion is conservative in that, for every L-sentence σ, $T'_f \models \sigma$ iff $T_f \models \sigma$.

Let T be a strongly pseudofinite (possibly incomplete) L-theory. Take T' to be the L'-consequences of $T \cup T'_f$.

We show that that T' is consistent. If otherwise, compactness yields $\sigma \in T$ inconsistent with T'_f, i.e. $T'_f \models \neg\sigma$. However, $\neg\sigma$ is an L-sentence, so $T_f \models \neg\sigma$. This conclusion contradicts even the weak psueofiniteness of T, which asserts that T (and therefore σ) is consistent with T_f.

Furthermore, we show that the expansion from T to T' is conservative in the same sense as the expansion from T_f to T'_f. It is sufficient to show that every L-sentence in T' is additionally in T. Let σ be such an L-sentence. Since $\sigma \in T'$, compactness yields that σ is a consequence of some $\tau \in T$ together with T'_f. Rephrasing, $T'_f \models \tau \to \sigma$. Since $\tau \to \sigma$ is an L-sentence, we know $T_f \models \tau \to \sigma$. Recall that strong psueofiniteness of T asserts that $T \models T_f$. Therefore, since $\tau \in T$, $\sigma \in T$.

We conceptualize T_f as containing all the first order principles that come from a structure's finiteness (e.g., all definable surjections from a set to itself are bijective). In Construction I.10, T'_f expands these principles to those involving counting and subsets. Naturally, for infinite models, this counting takes place in $\mathcal{R}^*$ rather than $\mathcal{R}$.

Exercise I.11. Use Construction I.10 to show that any saturated L-structure expands (not necessarily canonically) to an L'-structure.

From now on, whenever we have a pseudofinite structure or strongly pseudofinite theory, we will implicitly use symbols from L' and reason using the transfer principle (i.e., using sentences in T'_f). If our conclusion is a sentence in L, then that sentence holds via Construction I.10. We will use L to denote the original language, and L' to denote the expanded language.

Exercise I.12. Let M be a pseudofinite structure. Then a subset of M^n is definable in L' with parameters iff it corresponds to some set in $\mathcal{P}^*(M^n)$. Call these the *internal* subsets.

In addition to assigning nonstandard natural cardinalities to definable sets, we can define a real-valued measures on these sets.

Definition I.13. Let $X \subseteq M^n$ be definable without parameters. For any internal $Y \subseteq X$, define $\mu_X(Y)$ by

$$\mu_X(Y) = \mathrm{st}\left(\frac{|Y|}{|X|}\right)$$

Exercise I.14. Let M be a saturated, pseudofinite model, and let $X \subseteq M^n$ be L-definable without parameters. Show that μ_X is a finitely additive probability measure on the boolean algebra of internal subsets of X.

Exercise I.15. Let G be a definable group in a strongly pseudofinite theory. Show that $\mu_G(-)$ is both left and right invariant (i.e., for all definable $Y \subseteq G$, $g \in G$, we have $\mu_G(Y) = \mu_G(g \cdot Y) = \mu_G(Y \cdot g)$).

Hint: make an argument for finite groups and use the transfer principle.

The following lemma asserts that the measure of a definable set depends only on the L'-type of the parameter used in its definition.

Lemma I.16. *Let M be a saturated pseudofinite structure, let $X \subseteq M^n$ be L-definable without parameters, let $\bar{b}_0$ and $\bar{b}_1$ have the same L'-type p over $\emptyset$, and let Y_0 and Y_1 be subsets of X defined by $\varphi(\bar{y}, \bar{b}_0)$ and $\varphi(\bar{y}, \bar{b}_1)$, respectively. Then $|Y_0| = |Y_1|$.*

Proof. Note that, for any rational r in the unit interval, $\frac{|Y_0|}{|X|} > r$ iff $\frac{|Y_1|}{|X|} > r$, as decided by p. Take s to be the supremum of all the r such that those claims hold. It is clear that $\mu_X(Y_0) = \mu_X(Y_1) = s$. □

Remark I.17. More generally, for any φ and any closed $I \subseteq [0, 1]$,

$$\{\bar{b} \in M^m \mid \mu_X(\varphi(\bar{x}, \bar{b})) \in I\}$$

is type-definabe by L'-sentences over $\emptyset$.

Recall the following basic definition from stability theory.

Definition I.18. Let M be any saturated structure, $A \subseteq M$ be small, and $\varphi(\bar{x}, \bar{b})$ be a formula of L with parameters $\bar{b}$. We say that $\varphi(\bar{x}, \bar{b})$ *does not divide* over A if, for any infinite sequence $\bar{b} = \bar{b}_0, \bar{b}_1, \bar{b}_2, ...$ that is indiscernible over A, $\{\varphi(\bar{x}, \bar{b}_i) | i = 0, 1, ...\}$ is consistent with $\mathrm{Th}(M, a)_{a \in A}$, i.e., for all $k \in \omega$, $M \models \exists \bar{x}(\bigwedge_{i=1}^{k} \varphi(\bar{x}, \bar{b}_i))$.

Recall that a sequence of k-tuples $\bar{b}_0, \bar{b}_1, \bar{b}_2, ...$ is indiscernible over A iff, for all $i_0 < i_1, ... < i_n < \omega$ and for all $j_1 < j_2 < ... < j_k < \omega$, $\mathrm{tp}(\bar{b}_{i_0}, ..., \bar{b}_{i_n}/A) = \mathrm{tp}(\bar{b}_{j_0}, ..., \bar{b}_{j_n}/A)$. E.g.: suppose that $\bar{a}_1, \bar{a}_2, ..., \bar{a}_n$ is a sequence of n-tuples in M and that there is a set $\Sigma(\bar{x}_1, ..., \bar{x}_k)$ of formulas of L such that, for all $i_1 < ... < i_k$, $M \models \Sigma_k(\bar{a}_1, ..., \bar{a}_k)$. So, by Ramsey's theorem and Compactness, for all k, there is an indiscernible sequence $\bar{b}_1, ..., \bar{b}_k$ of tuples of M such that $M \models \Sigma_k(\bar{b}_1, ..., \bar{b}_k)$. [For a proof of this statement see Pillay, *Lecture notes . Model Theory*, Proposition 5.11, pp. 52-53.]

Lemma I.19. *Let M be a pseudofinite and saturated structure. Let $X \subseteq M^n$ be $\emptyset$-definable. For some $\bar{b} \in M^m$, let $\varphi(\bar{x}, \bar{b})$ define in L a subset of X. Suppose $\mu_X(\varphi(\bar{x}, \bar{b})(M)) > 0$. Then $\varphi(\bar{x}, \bar{b})$ does not divide over $\emptyset$ in the sense of L'.*

Proof. Let $\bar{b} = \bar{b}_0$ and let the sequence $\{\bar{b}_i\}_{i \in \omega}$ be indiscernible over $\emptyset$ in the sense of L'. We have to show that

$$\bigwedge_{i \in \omega} \varphi(\bar{x}, \bar{b}_i$$

is nonempty as a subset of M. It is sufficient to show that, for all n,

$$\mu_X\left(\bigwedge_{i<n}\varphi(\bar{x},\bar{b}_i)\right)>0.$$

Suppose otherwise, and take n least such that

$$\mu_X\left(\bigwedge_{i<n}\varphi(\bar{x},\bar{b}_i)\right)=0.$$

By assumption, since $\mu_X(\varphi(\bar{x},\bar{b}>0$, $n>1$. We therefore have

$$\mu_X\left(\bigwedge_{i<n-1}\varphi(\bar{x},\bar{b}_i\right)=a>0.$$

Define

$$Y_k=\bigwedge_{(n-1)k\leq i<(n-1)(k+1)}\varphi(\bar{x},\bar{b}_i),$$

and observe that, by indiscernibility and Lemma I.16, the various Y_k each have measure $a>0$ and are measure disjoint. This observation contradicts the finite additivity of μ_X. □

We finish this chapter with a few more examples of pseudofinite structures.

Example I.20. let L be a language with two sorts P and Q and one binary relation $\in$ between P and Q. Let $\mathcal{K}$ be the class of L-structures of the form $(X,\mathcal{P}(X),\in)$ where X is a finite set. Let $T_{\mathcal{K}}=Th(\mathcal{K})$ and let M be a saturated model of $T_{\mathcal{K}}$. M is pseudofinite. Furthermore, note that if $Y\subseteq P$ is definable with parameters in M, then there is some $a\in Q$ such that $Y=\{y\in P:M\models y\in a\}$.

Exercise I.21. Show that the theory of the random graph is pseudofinite. Hint: axiomatize the theory of random graph and use a probabilistic argument to show that, for any finite collection of axioms, there is a sufficiently large finite graph realizing those axioms (without explicitly constructing such a graph).

Chapter II

Graph Regularity Lemmas

In this section, we give some background on the classical Szemeredi Regularity Lemma, and prove a non-standard version by pseudofinite methods which implies the classical version.

Definition II.1. Let $G = (V, E)$ be a finite graph. The *density* of G is

$$d(G) = \frac{|E|}{|V|^2}.$$

For $A, B \subseteq V$, we say that

$$d(A, B) = \frac{|(A \times B) \cap E|}{|A||B|}.$$

We are going to study *regularity* of a graph $G = (V, E)$. Roughly speaking, *regularity* of (V, E) means: for all $A, B \subseteq V$, $d(G) \sim d(A, B)$(where $\sim$ means: approximately equal).

Definition II.2. Let $\epsilon > 0$. A graph (V, E) ϵ-*regular* if, for all $A, B \subseteq V$ such that $|A|, |B| \geq \epsilon|V|$, $|d(A, B) - d(G)| < \epsilon$.

Proposition II.3 (Szemeredi Regularity Lemma)**.** *Fix $\epsilon > 0$. There is $N(\epsilon) \in \mathcal{N}$ such that, for all finite graphs $G = (V, E)$, for some $N \leq N(\epsilon)$, there is a partition of the vertex set $V = V_1 \sqcup ... \sqcup V_N$ such that, for all $(i, j) \in \{1, ..., N\}^2$ (except at most ϵN^2 such pairs) and, for all $A \subseteq V_i, B \subseteq V_j$ such that $|A| \geq \epsilon|V_i|, |B| \geq \epsilon|V_j|$, $|d(A, B) - d(V_i, V_j)| \leq \epsilon$. Moreover, we can take our partition such that, for all i, j, $||V_i| - |V_j|| \leq 1$*

For $A, B \subseteq V$, we define $E(A, B) := |\{(a, b) \in A \times B : E(a, b)\}|$.

Proposition II.4 (Weak variant of II.3)**.** *For all ϵ, there is an $N(\epsilon)$ such that, for all finite $G = (V, E)$, for some $N \leq N(\epsilon)$, there is a partition $V = V_1 \sqcup ... \sqcup V_M$ and there is $S \subset \{(i, j) : 1 \leq i, j \leq N\}$ such that, for all $(i, j) \notin S$, for some $d_{ij} \in \mathcal{R}^{\geq 0}$, $|E(A, B) - d_{ij}|A||B|| \leq \epsilon|V_i||V_j|$ where $\sum_{(i,j) \in S} |V_i||V_j| \leq \epsilon|V^2|$.*

Exercise II.5. Prove Proposition II.4 from Proposition II.3.

Proposition II.6 (Non-Standard Version of II.4)**.** *For any pseudofinite graph G and any standard $\epsilon > 0$, there is a partition of the vertex set V into finitely many internal sets $V = V_1 \sqcup \ldots \sqcup V_N$. Furthermore, there is $S \subseteq \{(i,j) : 1 \leq i,j \leq N\}$ with $\sum_{(i,j)\in S} |V_i||V_j| < \epsilon |V|^2$ and, for all $(i,j) \notin S$, there is a standard $d_{ij} \geq 0$ such that, for all internal $A \subseteq V_i$ and $B \subseteq V_j$,*

$$|E(A,B) - d_{ij}|A||B|| \leq \epsilon\, |V_i|\, |V_j|\, .$$

We prove Proposition II.6 later in these notes (II).

Lemma II.7. *Proposition II.6 implies Proposition II.4.*

Proof. Suppose that II.4 fails. So, there is a standard rational $\epsilon > 0$ such that no N works. Take such an ϵ, and, for each number N, take G_N witnessing the failure N, and note that this failure is expressible as an L'- sentence. Take G to be a nonprinciple ultraproduct of the G_N graphs, and observe that G and ϵ witness the failure of Proposition II.6. □

The goal now is to give a proof of II.6. The proof will go via integration and measure theory. We will extend the finitely additive measures μ_X described above to actual σ-additive probability measures:

Proposition II.8. *Let M be a κ-saturated model, for κ a sufficiently large cardinal. Let $\mathcal{B}_0$ be the boolean algebra of definable (with parameters) subsets of M, and let $\mathcal{B}$ be the sigma algebra generated by $\mathcal{B}_0$. Let μ be a Kiesler measure on M, i.e. a finitely additive probability measure defined on $\mathcal{B}_0$. Then μ extends uniquely to a σ-additive probability measure defined on $\mathcal{B}$.*

We will give two proofs of this statement. The first, which follows now, is a straightforward application of the Carathéodory Extension Theorem. The second relies on type spaces and gives a bit more information.

Proof 1 of Proposition II.8. The Carathéodory Extension Theorem says that, whenever μ is a σ-additive measure defined on a boolean algebra $\mathbb{B}_0$, μ extends to the σ-algebra generated by $\mathbb{B}_0$. Furthermore, if μ is σ-finite, this extension is unique. We therefore need only check that μ is already σ-additive. (As a probability measure, μ is automatically σ-finite.)

Let X be defined by $\varphi(x)$, with parameters, and let

$$X = \bigsqcup_{i<\omega} X_i,$$

where each X_i is defined with parameters by $\varphi_i(x)$. We want to show that

$$\mu(X) = \sum_{i<\omega} \mu(X_i).$$

First, observe that

$$\varphi(x) \to \bigvee_{i<\omega} \varphi_i(x);$$

if it were otherwise, saturation of M would yield a point in $X \setminus \bigcup X_i$. Next, compactness yields an $n < \omega$ such that

$$\varphi(x) \to \bigvee_{i<n} \varphi_i(x),$$

so in particular

$$X = \bigcup_{i<n} X_i$$

and, for $i \geq n$, X_i is empty. The desired conclusion then follows from the finite additivity of μ. $\square$

Here are some unrelated exercises for you, dear reader.

Exercise II.9. Let L be a countable language and let $\mathcal{M}$ be a countable and $\aleph_0$-categorical L-structure. $\mathcal{M}$ (or $\mathrm{Th}(\mathcal{M})$ is said to be *smoothly approximable* if there is an ascending chain of finite substructures $A_0 \subseteq A_1 \subseteq \ldots \mathcal{M}$ such that:

1. $\bigcup_{i\in\omega} A_i = \mathcal{M}$ and,
2. for every i, and for every $\bar{a}, \bar{b} \in A_i$ if $\mathrm{tp}_{\mathcal{M}}(\bar{a}) = \mathrm{tp}_{\mathcal{M}}(\bar{b})$, then there is an automorphism σ of $\mathcal{M}$ such that $\sigma(\bar{a}) = \bar{b}$ and $\sigma(A_i) = A_i$.

Show that if $\mathcal{M}$ is smoothly approximable, then $\mathcal{M}$ is pseudofinite.

Exercise II.10. In Exercise I.21, you showed that the random graph is pseudofinite. Show that the random graph is not smoothly approximable.

Our aim here is to give a second proof of II.8 which provides more information than the first. For a more detailed exposition see Simon's notes on NIP theories [8]. We begin with some background on type spaces.

Let M be a saturated model (i.e., M is κ-saturated where $\kappa = |M|$), and let $M_0 \prec M$ be an elementary substructure of M_0 of size less than κ. The subsets of M definable with parameters from M_0 form a Boolean algebra which we call $\mathcal{B}_{M_0}$. It is easy to see that $\mathcal{B}_{M_0}$ is isomorphic to the Boolean algebra of formulas $\varphi(x)$ with parameters in M_0 up to equivalence modulo $\mathrm{Th}(M, a)_{a\in M_0}$, so we may identify these two freely in the discussion that follows.

Suppose we are given a Keisler measure μ on $\mathcal{B}_{M_0}$. Instead of working with subsets of M, it will be useful at first to reinterpret μ as a finitely additive probability measure on the collection of clopen subsets of the topological space $S_1(M_0)$ of complete 1-types over M_0, and extend this reinterpretation of μ to a full Borel probability measure on $S_1(M_0)$. These ideas are justified below.

We define the topology on $S_1(M_0)$ by declaring that $X \subseteq S_1(M_0)$ is closed if there is a set $\Sigma(x)$ of $\mathcal{L}_{M_0}$ formulas such that $X = \{p(x) | \Sigma(x) \subseteq p(x)\}$. Given a set $\Sigma(x)$ of $\mathcal{L}_{M_0}$ formulas, by $[\Sigma(x)]$ we mean the set

$$\{p(x) \in S_1(M_0) | \Sigma \subseteq p\}.$$

Note. $O \subseteq S_1(M_0)$ is open iff there is some collection Σ of $\mathcal{L}_{M_0}$-formulas with $O = \bigcup_{\varphi \in \Sigma} [\varphi]$

Proof. Clear from the definition of the topology on $S_1(M_0)$. $\square$

Exercise II.11. $S_1(M_0)$ is a compact, Hausdorff, totally disconnected space.

Proposition II.12. *$\mathcal{B}_{M_0}$ is isomorphic to the Boolean algebra $\mathcal{B}$ of clopen subsets of $S_1(M_0)$.*

Proof. We claim that the map $\varphi(M) \mapsto [\varphi(x)]$ is the desired isomorphism. Given subsets $\varphi_1(M)$ and $\varphi_2(M)$ of M defined by $\mathcal{L}_{M_0}$-formulas φ_1, φ_2, respectively, to show that the map is well-defined and preserves the Boolean algebra structure, it suffices to prove that $\varphi_1(M) \subseteq \varphi_2(M) \Leftrightarrow [\varphi_1(x)] \subseteq [\varphi_2(x)]$. Suppose $\varphi_1(M) \subseteq \varphi_2(M)$. If $p(x) \in [\varphi_1(x)]$, then (by saturation) $p(x) = \mathrm{tp}_M(a/M_0)$ for some $a \in M$, and $M \models \varphi_1(a) \Rightarrow a \in \varphi_1(M) \subseteq \varphi_2(M) \Rightarrow M \models \varphi_2(a) \Rightarrow p = tp_M(a/M_0) \in [\varphi_2(x)]$. To prove that the map is injective, note that if $[\varphi_1(x)] = [\varphi_2(x)]$, then $a \in \varphi_1(M) \Leftrightarrow \varphi_1(x) \in \mathrm{tp}_M(a/M_0) \Leftrightarrow \varphi_2(x) \in \mathrm{tp}_M(a/M_0) \Leftrightarrow a \in \varphi_2(M)$. Finally, surjectivity is clear. $\square$

Proposition II.12 is what allows us to work for the moment with $S_1(M_0)$ instead of with M, reinterpreting the given Keisler measure μ on $\mathcal{B}_{M_0}$ as a finitely additive probability measure defined on the Boolean algebra $\mathcal{B}$ of clopen subsets of $S_1(M_0)$. So for example, we interpret $\mu([x = x]) = \mu(M_0) = 1$, and in general, $\mu([\varphi(\bar{x})]) = \mu(\varphi(M_0))$.

Proposition II.13. *μ entends to a regular Borel probability measure on $S_1(M_0)$.*

Recall that a Borel probability measure μ on a compact space X is *regular* if for any Borel set $B \subseteq X$ and $\epsilon > 0$ there is a closed set $C \subset B$ and an open set $U \supset B$ with $\mu(U) - \mu(C) < \epsilon$.

Remark II.14. A regular Borel probability measure μ on $S_1(M_0)$ is uniquely determined by its restriction to the clopen subsets of $S_1(M_0)$. To see why, suppose $C \subseteq S_1(M_0)$ is closed. By regularity, $\mu(C) = \inf\{\mu(O) | O \supset C \text{ open}\}$. But given any open set $O \supset C$, since O is a union $\bigcup D_i$ of clopen sets, and C is compact, there is a clopen set D such that $C \subset D \subset O$. Thus $\mu(C) = \inf\{\mu(D) | D \supset C \text{ is clopen}\}$ is determined by the measures of clopens. Now suppose X is any Borel subset of $S_1(M_0)$. Again, by regularity of μ, we have $\mu(X) = \sup\{\mu(C) | C \subset X \text{ is closed}\}$. Thus $\mu(X)$ is determined by the measures of the closed subsets of X, which in turn are determined by measures of clopens.

This type of uniqueness is not exactly what we will need in the end. Note that the σ-algebra $\sigma\mathcal{B}$ generated by the clopens in $S_1(M_0)$ may be a proper sub σ-algebra of the Borel σ-algebra $\sigma\widetilde{\mathcal{B}}$ (there could be open sets in $S_1(M_0)$ not given by any countable collection of formulas). Further, while the probability measure we will define on $\sigma\widetilde{\mathcal{B}}$ is the unique regular extension of μ to the Borels, it need not be unique among all (possibly non regular) probability measures extending μ to $\sigma\widetilde{\mathcal{B}}$. However, none of these concerns is a problem when we restrict our attention to $\sigma\mathcal{B}$.

Corollary II.15. *There is a unique extension of μ to a (σ-additive) probability measure on the σ-algebra $\sigma\mathcal{B}$ generated by the clopen subsets of $S_1(M_0)$.*

Proof. II.13 gives the existence of an extension $\overline{\mu} : \sigma\mathcal{B} \to [0, 1]$. For uniqueness, it is enough to show that this (and every other extension) coincides with the outer measure generated by μ. But this is a standard argument that can be found in any elementary real analysis textbook. $\square$

Definition II.16. For any subset X of $S_1(M_0)$, we say X satisfies the *regularity property with respect to* μ if

$$\sup\{\mu(F)|F \subset X \text{ closed}\} = \inf\{\mu(O)|X \subset O \text{ open}\}$$

and we define $\mu(X)$ to be $\sup\{\mu(F)|F \subset X \text{ closed}\}$.

Note that, since $S_1(M_0)$ is compact (and hence, every closed subset is compact), saying that μ extends to a measure satisfying the regularity property for every Borel subset is equivalent to saying that μ extends to a regular Borel measure on $S_1(M_0)$.

Proof of II.13. The goal is to extend the measure μ so that the regularity property with respect to μ holds for all Borel subsets of $S_1(M_0)$ i.e. to extend μ to a regular Borel measure on $S_1(M_0)$.

For $O \subset S_1(M_0)$ open, define

$$\mu(O) := \sup\{\mu(D)|D \subset O \text{ clopen}\},$$

and for $C \subset S_1(M_0)$ closed, let

$$\mu(C) := \inf\{\mu(D)|D \supset C \text{ clopen}\}.$$

We claim that for any closed (and similarly, open) $X \subset S_1(M_0)$, X satisfies the regularity property with respect to μ. Suppose X is closed. Then

$$\begin{aligned}\inf\{\mu(O)|X \subset O \text{ open}\} &= \inf\{\mu(D)|X \subset D \text{ clopen}\} \quad \text{(Since X is compact)}\\ &= \mu(X)\\ &= \sup\{\mu(F)|F \subset X \text{ closed}\}\end{aligned}$$

and similarly for X open.

For arbitrary subsets $X \subseteq S_1(M_0)$, if X satisfies regularity with respect to μ, we define

$$\mu(X) := \sup\{\mu(F) | F \subset X \text{ closed}\} = \inf\{\mu(O) | X \subset O \text{ open}\}.$$

Since μ is regular on open and closed subsets of $S_1(M_0)$, to see that μ extends to a regular measure on the Borel subsets, it suffices to show that the collection of sets satisfying the regularity property is a σ-algebra and that μ is countably additive of this σ-algebra. By definition, clopen sets satisfy the regularity property with respect to μ, so we have $\mu(\emptyset) = 0$ and $\mu([x = x]) = 1$. Clearly if $A \subseteq B$ are sets satisfying the regularity property with respect to μ, then $\mu(A) \leq \mu(B)$, since any open $O \subset A$ is also an open subset of B.

To see that the collection of sets satisfying regularity is closed under complementation, suppose X satisfies the regularity property, then

$$\begin{aligned}\inf\{\mu(O) | X^C \subset O \text{ open}\} &= \inf\{1 - \mu(F) | F \subset X \text{ closed}\} \\ &= 1 - \sup\{\mu(F) | F \subset X \text{ closed}\} \\ &= 1 - \inf\{\mu(O) | X \subset O \text{ open}\} \\ &= \sup\{1 - \mu(O) | X \subset O \text{ open}\} \\ &= \sup\{\mu(F) | F \subset X^C \text{ closed}\}.\end{aligned}$$

and hence X^c satisfies the regularity property.

For countable union, we first prove a small claim: suppose that if $X_1, \ldots, X_n$ are pairwise disjoint sets satisfying the regularity property with respect to μ, then $X := X_1 \sqcup \ldots \sqcup X_n$ also satisfies the regularity property and $\mu(X) = \sum_i \mu(X_i)$. Clearly it suffices to prove the claim for $n = 2$ and proceed by induction. To show this, we first show two special cases:

1. If $X_i = F_i$ are closed and $X = F = F_1 \sqcup F_2$, then F is closed (and hence satisfies regularity) and it is clear from the definition that $\mu(F) \leq \mu(F_1) + \mu(F_2)$. For the reverse inequality, fix $\epsilon > 0$ and let $D \supseteq F$ be a clopen set such that $\mu(D) - \mu(F) < \epsilon$. Further, choose $D_1 \supseteq F_1$, and $D_2 \supseteq F_2$ such that $\mu(D_i) - \mu(F_i) < \epsilon$ for $i = 1, 2$. Let $D_i' := D_i \cap D$. Then

$$\begin{aligned}\mu(F_1) + \mu(F_2) &\leq \mu(D_1') + \mu(D_2') \\ &= \mu(D \cap (D_1 \sqcup D_2)) \quad \text{(since } \mu \text{ is finitely additive for clopen sets)} \\ &\leq \mu(D) \\ &< \mu(F) + \epsilon.\end{aligned}$$

2. If $X_i = O_i$ are open and $X = O = O_1 \sqcup O_2$, then, again, O is open and so satisfies regularity, and the inequality $\mu(O) \geq \mu(O_1) + \mu(O_2)$ is clear from the definition of μ on open sets. For the reverse, fix $\epsilon > 0$ and a clopen set $D \subseteq O$ satisfying $\mu(O) - \mu(D) < \epsilon$. Then

$$D \subseteq O_1 \sqcup O_2 = \left(\bigcup_{i \in I} [\varphi_i] \right) \sqcup \left(\bigcup_{j \in J} [\psi_j] \right)$$

for some collections $\{\varphi_i\}_{i\in I}$ and $\{\psi_j\}_{j\in J}$ of formulas. Since D is compact, we can find clopen sets

$$D_1 = \left[\bigvee_{1\le k\le n} \varphi_{i_k}\right] \subset O_1 \text{ and } D_2 = \left[\bigvee_{1\le k\le m} \psi_{j_k}\right] \subset O_2$$

where $D \subseteq D_1 \sqcup D_2$. Then

$$\begin{aligned}\mu(O) &< \mu(D) + \epsilon \\ &\le \mu(D_1 \sqcup D_2) + \epsilon \\ &= \mu(D_1) + \mu(D_2) + \epsilon \\ &\le \mu(O_1) + \mu(O_2) + \epsilon.\end{aligned}$$

Now for general X_1, X_2 satisfying the regularity property, fix $\epsilon > 0$, and for $i = 1, 2$, choose closed sets F_i and open sets O_i with $F_i \subset X_i \subset O_i$ so that $\mu(O_i) - \mu(F_i) < \epsilon$. Then $F := F_1 \sqcup F_2$ and $O := O_1 \cup O_2$ satisfy $F \subset X \subset O$, and

$$\begin{aligned}\mu(O) - \mu(F) &\le \mu(O_1) + \mu(O_2) - (\mu(F_1) + \mu(F_2)) \\ &= \mu(O_1) - \mu(F_1) + \mu(O_2) - \mu(F_2) \\ &< 2\epsilon\end{aligned}$$

and hence X satisfies regularity. The inequality $\mu(X) \ge \mu(X_1) + \mu(X_2)$ follows in a straightforward manner from the definition of μ.
For the reverse inequality, fix $\epsilon > 0$. Let $O_i \supset X_i$ with $\mu(O_i) - \mu(X_i) < \epsilon$. Then

$$\begin{aligned}\mu(X) &\le \mu(O_1 \cup O_2) \\ &\le \mu(O_1) + \mu(O_2) \\ &\le \mu(X_1) + \mu(X_2) + 2\epsilon.\end{aligned}$$

This proves the claim. Observe that the claim gives us finite subadditivity (an easy exercise in measure theory).

Now suppose $X = \cup_{i\in\omega} X_i$ and the regularity property holds for all X_i. Fix $\epsilon > 0$. Since each X_i satisfies the regularity property we may choose, for each i, sets $F_i \subset X_i \subset O_i$, with F_i closed and O_i open, such that $\mu(O_i) - \mu(F_i) < \frac{\epsilon}{2^i}$. Let $O := \bigcup_{i\in\omega} O_i$. Since O is open,

$$\mu(O) = \sup\left\{\mu(K) : K \text{ clopen, } (\exists n \in \omega) K \subseteq \bigcup_{i=0}^{n} O_i\right\}$$

and so there is some $N \in \omega$ such that for $\tilde{O} := \cup_{i\le N} O_i$ we have

$$\mu(O) - \mu(\tilde{O}) \le \mu(O) - \mu(D) < \epsilon.$$

Let $F := \cup_{i \leq N} F_i$. Then $F \subset X \subset O$, and

$$\begin{aligned}\mu(O) - \mu(F) &= \mu(O) - \mu(\tilde{O}) + \mu(\tilde{O}) - \mu(F) \\ &< \epsilon + \mu(\tilde{O} \setminus F).\end{aligned}$$

Observe now that, by the previous claim,

$$\begin{aligned}\mu(\tilde{O}) &= \mu(\tilde{O} \setminus F \sqcup F) \\ &= \mu(\tilde{O} \setminus F) + \mu(F)\end{aligned}$$

since $\tilde{O} \setminus F$ is open and so satisfies regularity and similarly for F since F is closed. Therefore $\mu(\tilde{O} \setminus F) = \mu(\tilde{O}) - \mu(F)$. Now, we have

$$\begin{aligned}\mu(O) - \mu(F) &< \epsilon + \mu(\tilde{O} \setminus F) \\ &= \epsilon + \mu(\tilde{O}) - \mu(F) \\ &= \epsilon + \mu(\bigcup_{i \leq N} O_i) - \mu(\bigcup_{i \leq N} F_i).\end{aligned}$$

By finite subadditivity for sets satisfying regularity,

$$\begin{aligned}\mu(O) - \mu(F) &\leq \epsilon + \sum_{i \leq N} \mu(O_i) - \sum_{i \leq N} \mu(F_i) \\ &= \epsilon + \sum_{i \leq N} \mu(O_i) - \mu(F) \\ &< \epsilon + \sum_{i \leq N} \frac{\epsilon}{2^i} \\ &< 3\epsilon.\end{aligned}$$

Since ϵ was arbitrary, $\sup\{\mu(F) : F \subset X \text{ closed}\} = \inf\{\mu(O) : X \subset O \text{ open}\}$ and hence X satisfies regularity with respect to μ, as required. Thus, we have showed that the sets satisfying the regularity property form a σ-algebra contains the σ-algebra of Borel sets.

Finally, we show that μ is indeed additive on sets satisfying the regularity property, and so is a regular Borel probability measure on $S_1(M_0)$. Let $X = \bigsqcup_{i \in \omega} X_i$ with each X_i satisfying the regularity property. Given $n \in \omega$, let $\widetilde{X}_n = \bigcup_{i \leq n} X_i$. Then $\mu(\widetilde{X}_n) = \sum_{i \leq n} \mu(X_i)$, and $\mu(\widetilde{X}_n) \leq \mu(X)$ since $\widetilde{X}_n \subseteq X$. Since the sequence $\{\mu(\widetilde{X}_n)\}_{n \in \omega}$ is increasing and bounded above by $\mu(X)$, it follows that the series $\sum_{i \leq n} \mu(X_i)$ converges to a value at most $\mu(X)$. To see that $\mu(X) \leq \sum_{i < \omega} \mu(X_i)$, fix $\epsilon > 0$ and choose $N \in \omega$ large enough so that $\sum_{i > N} \mu(X_i) < \epsilon$. Then $X = \widetilde{X}_N \sqcup (\bigcup_{i > N} X_i)$, so

$$\mu(X) = \mu(\widetilde{X}_N) + \mu(\bigcup_{i > N} X_i) < \mu(\widetilde{X}_N) + \epsilon \leq \sum_{i \in \omega} \mu(X_i) + \epsilon.$$

$\square$

We have focused our attention on the type space $S_1(M_0)$. We now need to prove that the unique extension of μ to the σ-algebra $\sigma\mathcal{B}$ generated by the clopen subsets of $S_1(M_0)$ can be interpreted as a well-defined (and unique) extension of the original Keisler measure μ on $\mathcal{B}_{M_0}$ to $\sigma\mathcal{B}_{M_0}$. To achieve this, it is enough to show that the σ-algebras $\sigma\mathcal{B}_{M_0}$ and $\sigma\mathcal{B}$ are isomorphic, or that there is a bijection between them preserving countable unions and complements.

Proposition II.17. *Fix collections $\mathcal{B}$, $\mathcal{C}$ of subsets of M and $S_1(M_0)$, respectively, such that for every $x, y \in M$, and every $X \in \mathcal{B}$, if $\mathrm{tp}_M(x/M_0) = \mathrm{tp}_M(y/M_0)$, then $x \in X \Leftrightarrow y \in X$.*
Suppose further that the map $f : \mathcal{B} \to \mathcal{C}$ given by $f(x) = \{\mathrm{tp}_M(a/M_0) | a \in X\}$ is a well-defined bijection.
Then the map

$$\begin{aligned}\widetilde{f} : \sigma\mathcal{B} &\to \sigma\mathcal{C} \\ X &\mapsto \{\mathrm{tp}_M(a/M_0) | a \in X\}\end{aligned}$$

is an isomorphism of σ-algebras.

Proof. Let $\mathcal{B} =: \mathcal{B}_0$, $\mathcal{C} =: \mathcal{C}_0$, and $f =: f_0$. Given $\mathcal{B}_\alpha$, $\mathcal{C}_\alpha$, and f_α, first let

$$\mathcal{B}^{(0)}_{\alpha+1} = \left\{ \bigcup_{i<\omega} X_i \middle| \{X_i\}_{i\in\omega} \subset \mathcal{B}_\alpha \right\},$$

and

$$\mathcal{C}^{(0)}_{\alpha+1} = \left\{ \bigcup_{i<\omega} X_i \middle| \{X_i\}_{i\in\omega} \subset \mathcal{C}_\alpha \right\}.$$

Then take

$$\begin{aligned}\mathcal{B}_{\alpha+1} &= \mathcal{B}^{(0)}_{\alpha+1} \cup \{X^c | X \in \mathcal{B}^{(0)}_{\alpha+1}\}, \text{ and} \\ \mathcal{C}_{\alpha+1} &= \mathcal{C}^{(0)}_{\alpha+1} \cup \{X^c | X \in \mathcal{C}^{(0)}_{\alpha+1}\}.\end{aligned}$$

Finally, define

$$\begin{aligned}f_{\alpha+1} &: \mathcal{B}_{\alpha+1} \to \mathcal{C}_{\alpha+1} \qquad \text{by} \\ &: X \mapsto \{\mathrm{tp}_M(a/M_0) | a \in X\}.\end{aligned}$$

If α is a limit ordinal, then let

$$\begin{aligned}\mathcal{B}_\alpha &= \bigcup_{\gamma<\alpha} \mathcal{B}_\gamma, \\ \mathcal{C}_\alpha &= \bigcup_{\gamma<\alpha} \mathcal{C}_\gamma, \text{ and}\end{aligned}$$

$$f_\alpha = \bigcup_{\gamma<\alpha} f_\gamma.$$

It is not difficult to show by induction on ordinals that for every α, we have the following properties:

(I) $\mathcal{B}_\alpha \subseteq \sigma\mathcal{B}$, and $\mathcal{C}_\alpha \subseteq \sigma\mathcal{C}$;

(II) Given $X \in \mathcal{B}_\alpha$, and $x, y \in M$, with $\mathrm{tp}_M(x/M_0) = \mathrm{tp}_M(y/M_0)$, then $x \in X$ iff $y \in X$;

(III) f_α is a well-defined bijection;

(IV) Given $X_i \in \mathcal{B}_{\gamma_i}$, with $i \in \omega$, and $\gamma_i < \alpha$ for every i, we have

$$f_\alpha\left(\bigcup_{i<\omega} X_i\right) = \bigcup_{i<\omega} f(X_i);$$

(V) Given $X \in \mathcal{B}_\gamma$ with $\gamma < \alpha$, we have $f_\alpha(X^c) = (f_\alpha(X))^c$.

Certainly by stage $\alpha = \omega_1$ we have an isomorphism of σ-algebras. For instance, given

$$\{X_i\}_{i<\omega} \subset \mathcal{B}_{\omega_1},$$

fix for each i some $\gamma_i < \omega_1$ with $X_i \in \mathcal{B}_{\gamma_i}$. Then $\alpha := \bigcup_{i<\omega} \gamma_i$ is countable, so $\alpha < \omega_1$, and thus $\bigcup_{i\in\omega} X_i \in \mathcal{B}_{\alpha+1} \subset \mathcal{B}_{\omega_1}$.

Finally, by property (I) from above, $\mathcal{B}_{\omega_1} \subset \sigma\mathcal{B}$ and $\mathcal{C}_{\omega_1} \subset \sigma\mathcal{C}$. This concludes the proof. □

Corollary II.18. *Let $\mathcal{B}$ be the collection of subsets of M defined over M_0, and $\mathcal{C}$ the collection of clopen subsets of $S_1(M_0)$. Then $\sigma\mathcal{B}$ and $\sigma\mathcal{C}$ are isomorphic as σ-algebras. In particular, the unique extension of $\mu : \mathcal{C} \to [0,1]$ to a probability measure on $\sigma\mathcal{C}$ corresponds to a unique extension of $\mu : \mathcal{B} \to [0,1]$ to a probability measure on $\sigma\mathcal{B}$.*

We now give a second proof of II.8, in which the Boolean algebra of definable sets consists of sets definable with parameters from all of M, not just with parameters from some small elementary substructure $M_0 \prec M$. First, we need a lemma.

Lemma II.19. *Given a Boolean algebra $\mathcal{B}$ of sets, any element of the σ-algebra $\sigma\mathcal{B}$ generated by $\mathcal{B}$ is contained in the σ-algebra generated by some countable subalgebra $\mathcal{B}' \subset \mathcal{B}$.*

Proof. Construct $\sigma\mathcal{B}$ in stages in the same way as in the previous proposition, so that $\sigma\mathcal{B} = \mathcal{B}_{\omega_1}$. Prove, by induction on α, that given $X \in \mathcal{B}_\alpha$, there exists a countable sequence $\{X_i\}_{i\in\omega} \subset \mathcal{B}$ such that X is in the σ-algebra generated by $\{X_i\}_{i\in\omega}$. □

Theorem II.8. *Let M be a large saturated model, and μ a Keisler measure on the Boolean algebra $\mathcal{B}$ of subsets of M definable with parameters from M. Then μ extends uniquely to a σ-additive probability measure μ' on the σ-algebra $\sigma\mathcal{B}$ generated by $\mathcal{B}$.*

Proof 2 of II.8. Fix $X \in \sigma\mathcal{B}$. Let $M_0 \prec M$ be a small elementary substructure of M such that $X \in \sigma\mathcal{B}_0$, where $\mathcal{B}_0$ is the Boolean algebra of subsets of M definable with parameters from M_0. (The previous lemma guarantees that such an elementary substructure exists.) Let μ_0 be the unique probability measure extending $\mu_{|\mathcal{B}_0}$ to all of $\sigma\mathcal{B}_0$. Define $\mu'(X) := \mu_0(X)$. We claim that $\mu'(X)$ does not depend on the choice of M_0: Fix small substructures $M_0 \prec M$ and $M_1 \prec M$ with $X \in \sigma\mathcal{B}_0 \cap \sigma\mathcal{B}_1$, where $\mathcal{B}_i$ is the Boolean algebra of sets definable over M_i. Let $M_2 \prec M$ be a small elementary substructure containing both M_0 and M_1, and let $\mathcal{B}_2$ be the Boolean algebra of sets definable over M_2. For $i = 0, \dots, 2$, denote by $\mu_i : \sigma\mathcal{B}_i \to [0,1]$ the unique extension of $\mu_{|\mathcal{B}_i}$ to all of $\sigma\mathcal{B}_i$. Since $\sigma\mathcal{B}_i \subset \sigma\mathcal{B}_2$ for $i = 0, 1$, we have $\mu_0(X) = \mu_2(X) = \mu_1(X)$, and hence μ' does not depend on the choice of M_0, as claimed.

Finally, to see that μ' is indeed a probability measure on $\sigma\mathcal{B}$, note that given any countable sequence $\{X_i\}_{i\in\omega} \subset \sigma\mathcal{B}$, there is some small elementary substructure $M_0 \prec M$ such that for every i, $X_i \in \sigma\mathcal{B}_0$. This concludes the proof of II.8. □

Recall now the context in which we originally planned to use II.8. We had a saturated pseudofinite graph $M = (V, E)$ with all the additional properties defined in Construction I.13 (I.9,I.10), and a Keisler measure μ_V on the definable subsets of V. We use the extension given by II.8 and some measure theoretic techniques to prove II.6 (the nonstandard version of Weak Szemeredi's Theorem). First, we present a brief review of integration on measure spaces.

Definition II.20. We consider a probability measure space $\{\Omega, \mathcal{B}, \mu\}$, where Ω is a set, $\mathcal{B}$ is a σ-algebra of subsets of Ω, and $\mu : \mathcal{B} \to [0,1]$ is a probability measure. Fix a function $f : \Omega \to \mathbb{R}_{\geq 0}$.

1. f is μ-measurable if the pre image $f^{-1}(B)$ of every Borel set $B \subset \mathbb{R}_{\geq 0}$ is an element of $\mathcal{B}$.

2. f is simple if it is μ-measurable and its range is finite, i.e $Range(f) = \{a_1, \dots, a_n\}$.

3. If f satisfies (II), then we define $\int\limits_\Omega f \ d\mu := \sum\limits_{i=1}^{n} a_i \ \mu(E_i)$, where $E_i = f^{-1}(\{a_i\})$.

4. For f as in (II), if $E \in \mathcal{B}$, we define $\int\limits_E f \ d\mu := \sum\limits_{i=1}^{n} a_i \ \mu(E_i \cap E)$.

In general, we can show that for every measurable, nonnegative function f, there is a sequence $\{f_i\}_{i\in\omega}$ of simple functions such that $f = \sup\limits_n f_n$. We then define $\int\limits_\Omega f \ d\mu := \sup\limits_n \int\limits_\Omega f_n \ d\mu$, and $\int\limits_E f \ d\mu := \sup\limits_n \int\limits_E f_n \ d\mu$, for any $E \in \mathcal{B}$.

Fact II.21 (Conditional Expectation)**.** *Consider a measure space* $\{\Omega, \mathcal{B}, \mu\}$*, with* μ *σ-finite, and let* $\mathcal{A}$ *be a σ-subalgebra of* $\mathcal{B}$*. Given a* $\mathcal{B}$*-measurable function* $f : \Omega \to \mathbb{R}_{\geq 0}$*, there is a unique (up to differences on sets of measure zero)* $\mathcal{A}$*-measurable function* $g : \Omega \to \mathbb{R}_{\geq 0}$ *with the property that* $\int_E g \; d\mu = \int_E f \; d\mu$ *for every* $E \in \mathcal{A}$*.* g *is denoted by* $\mathbb{E}(f|\mathcal{A})$*.*

Proof. This is an application of the Radon-Nikodym theorem. □

Exercise II.22. Let $\{\Omega, \mathcal{B}, \mu\}$ be a measure space, and fix A, $B \in \mathcal{B}$ where $\mu(B) \neq 0, 1$. Then $\mathcal{A} = \{\emptyset, B, B^c, \Omega\}$ is a σ-subalgebra of $\mathcal{B}$. Compute $\mathbb{E}(A|\mathcal{A})$, where $\mathbb{E}(A|\mathcal{A}) := \mathbb{E}(\chi_A|\mathcal{A})$.

We now restate and prove Proposition II.6:

Theorem II.6. *Let* $G = (V, E)$ *be a saturated pseudofinite graph (in the sense of Example I.21). Fix* $\epsilon > 0$*. Then there exists a standard natural number* m*, and a partition* $V = V_1 \cup \cdots \cup V_m$ *of* V *into* m *disjoint definable sets such that for every pair* (i, j) *not in an exceptional set* S *of pairs satisfying*

$$\sum_{(i,j)\in S} |V_i||V_j| \leq \epsilon |V|^2,$$

there exists a constant $d_{ij} \in \mathbb{R}_{\geq 0}$ *such that for all definable subsets* $A \subset V_i$ *and* $B \subset V_j$*,*

$$|E(A, B) - d_{ij}|A||B|| \leq \epsilon |V_i||V_j|.$$

Proof of II.6. Let $\mathcal{B}_V$ be the σ-algebra generated by the definable subsets of V, and μ_V the unique extension of the Keisler measure given by $X \mapsto \mathrm{st}\left(\frac{|X|}{|V|}\right)$ to $\mathcal{B}_V$. Let $\mathcal{B}_{V\times V}$ denote the σ-algebra generated by definable subsets of $V \times V$, and $\mu_{V\times V}$ the unique extension of the Keisler measure given by $X \mapsto \mathrm{st}\left(\frac{|X|}{|V|^2}\right)$ to $\mathcal{B}_{V\times V}$. Finally, let $\mathcal{B}_V \times \mathcal{B}_V$ be the σ-algebra generated by products $A \times B$ where $A \in \mathcal{B}_V$ and $B \in \mathcal{B}_V$, and $\mu_V \times \mu_V$ the measure $\mu_{V\times V}|_{\mathcal{B}_V\times\mathcal{B}_V}$.
Note that in general, $\mathcal{B}_V \times \mathcal{B}_V$ is a proper σ-subalgebra of $\mathcal{B}_{V\times V}$, and in particular the edge relation E need not belong to $\mathcal{B}_V \times \mathcal{B}_V$. Thus the function χ_E is $\mathcal{B}_{V\times V}$-measurable, (since E is a definable subset of $V \times V$), but not necessarily $\mathcal{B}_V \times \mathcal{B}_V$-measurable.
Let $f := \mathbb{E}(\chi_E|\mathcal{B}_V \times \mathcal{B}_V)$. Then f is $\mathcal{B}_V \times \mathcal{B}_V$-measurable, and by Fact II.21, for $A, B \subset V$ definable,

$$\int_{A\times B} f \; d\mu_{V\times V} = \int_{A\times B} \chi_E \; d\mu_{V\times V} = \mu_{V\times V}(E \cap (A \times B)) = \mathrm{st}\left(\frac{|E \cap (A \times B)|}{|V|^2}\right).$$

In fact, f has values in $[0, 1]$, so $||f|| = \int f$. We work for the moment in $\mathcal{B}_V \times \mathcal{B}_V$. Fix $\epsilon > 0$. Recall that f can be approximated by a sequence of simple functions $\{g_i\}$ and $\int f$ can be approximated by the integrals of those functions $\{\int g_i\}$. Recall also the fact that sets in $\mathcal{B}_V \times \mathcal{B}_V$ are approximated in $\mu_V \times \mu_V$-measure from below by products $A \times B$ of definable subsets A, B of V , and so we can partition V into definable sets $V_1 \cup \cdots \cup V_n$ and find constants $d_{ij} \geq 0$ such that

$$\sum_{ij} d_{ij} \chi_{V_i \times V_j} \leq f$$

and

$$\int\limits_{V \times V} \left(f - \sum_{ij} d_{ij} \chi_{V_i \times V_j} \right) \, d_{\mu_V \times \mu_V} < \epsilon^2.$$

Now $f \geq d_{ij}$ on $V_i \times V_j$, hence

$$\sum_{ij} \left(\int\limits_{V_i \times V_j} (f - d_{ij}) \right) < \epsilon^2. \qquad (*)$$

We claim that

$$\int\limits_{V_i \times V_j} (f - d_{ij}) \, d\mu_{V \times V} < \epsilon \cdot \mu_V(V_i) \cdot \mu_V(V_j)$$

for all (i, j) outside of some exceptional set S such that

$$\sum_{(i,j) \in S} \mu_V(V_i) \mu_V(V_j) < \epsilon.$$

Indeed, let $S = \left\{ (i, j) | \int\limits_{V_i \times V_j} (f - d_{ij}) \geq \epsilon \mu(V_i) \mu(V_j) \right\}$. Suppose for a contradiction that

$$\sum_{(i,j) \in S} \mu_V(V_i) \mu_V(V_j) \geq \epsilon.$$

Then

$$\begin{aligned} \sum_{(i,j) \in S} \int\limits_{V_i \times V_j} f - d_{ij} &\geq \epsilon \sum_{(i,j) \in S} \mu_V(V_i) \mu_V(V_j) \\ &\geq \epsilon \cdot \epsilon = \epsilon^2. \end{aligned}$$

This contradicts $(*)$ from above. We now fix $(i, j) \notin S$, and $A \subset V_i$, $B \subset V_j$ definable subsets of V. Then

$$\int\limits_{A \times B} f \, d\mu_{V \times V} - \int\limits_{A \times B} d_{ij} \, d\mu_{V \times V} = \int\limits_{A \times B} (f - d_{ij}) \, d\mu_{V \times V} < \epsilon \, \mu_V(V_i) \cdot \mu_V(V_j),$$

where

- $\displaystyle\int_{A\times B} f\, d\mu_{V\times V} = \int_{A\times B} \chi_E\, d\mu_{_{V\times V}} = \mu_{_{V\times V}}(E\cap(A\times B) = \text{st}\left(\frac{|E(A,B)|}{|V|^2}\right)$,
- $\displaystyle\int_{A\times B} d_{ij}\, d\mu_{V\times V} = d_{ij}\cdot\mu_{V\times V}(A\times B) = d_{ij}\ \text{st}\left(\frac{|A|\cdot|B|}{|V|^2}\right)$,
- $\mu_V(V_i) = \text{st}\left(\frac{|V_i||}{|V|}\right)$,

and

- $\mu_V(V_j) = \text{st}\left(\frac{|V_j|}{|V|}\right)$.

Putting these together and applying the properties of the standard part map yields:

$$\text{st}\left(\frac{|E(A,B)|}{|V|^2}\right) - d_{ij}\ \text{st}\left(\frac{|A|\cdot|B|}{|V|^2}\right) < \epsilon\left(\text{st}\left(\frac{|V_i|}{|V|}\right)\text{st}\left(\frac{|V_j|}{|V|}\right)\right)$$

and so

$$\text{st}\left(\left|\frac{|E(A,B)|}{|V|^2} - d_{ij}\ \frac{|A||B|}{|V|^2}\right|\right) < \text{st}\left(\epsilon\cdot\frac{|V_i||V_j|}{|V|^2}\right),$$

which gives

$$\left|\frac{|E(A,B)|}{|V|^2} - d_{ij}\ \frac{|A||B|}{|V|^2}\right| < \epsilon\cdot\frac{|V_i||V_j|}{|V|^2}$$
$$\Leftrightarrow |E(A,B) - d_{ij}\ |A||B|| < \epsilon|V_i||V_j|.$$

as required. □

Chapter III

Algebraically Closed Fields

This chapter will serve as a preliminary for our study of pseudofinite fields. For the most part, we will leave algebraic facts without proof, and refer the reader to any standard text on algebra for the details. We will consider fields as structures in the language of unitary rings $\mathcal{L}_{ring} = \{+, x, -, 0, 1\}$. A field K is said to be algebraically closed if every non-constant polynomial $P(x) \in K[x]$ has a solution in K. We denote the theory of algebraically closed fields of characteristice p by ACF_p. By K^{alg} we mean the algebraic closure of an arbitrary field K. If F is a subfield of K, we write $F < K$. It is convenient to recall the following facts about algebraically closed fields and algebraic varieties:

Fact III.1. *Let ACF_p denote the theory of algebraically closed fields of characteristic p. Then:*

1. *ACF has quantifier elimination;*
2. *the completions of ACF are precisely ACF_0 and ACF_p for p prime.*

Note that, since the language of rings with unity has constant symbols, any sentence σ of ACF is equivalent to a quantifer-free sentence τ where τ is Boolean combination of sentences of the form $t_1 = t_2$, where t_i are terms of the language (i.e. polynomial combinations definable elements).

Corollary III.2. *Any algebraically closed field K is strongly minimal. That is, any definable (with parameters in K) subset of K is finite or cofinite and the same holds for any elementary extension K_1 of K.*

Proof. Since ACF eliminates quantifiers, $Th(K, a)_{a\in K}$ does as well. Thus, any quantifier-free formula $\varphi(x)$ (in one variable) is logically equivalent to a Boolean combination of terms of the form '$P(x) = 0$', where $P(x)$ is a polynomial over K. Either $P(x)$ is trivial, i.e the zero polynomial, or it is not. If $P(x)$ is not trivial, then $P(x)$ is polynomial of non-zero finite degree, and thus, has finitely many solutions. Therefore, the solution set for $P(x) = 0$ is either finite or is everything in the field (and, clearly, the solution set of $P(x) \neq 0$ is either empty or cofinite). The class of finite or cofinite subsets of K is closed under finite Boolean combinations, and so any definable subset of K is finite or cofinite. $\square$

Next, we give two different definitions of what it means for a field element to be *algebraic* over a field. In one instance, we will mean algebraic in the sense of classical field theory. In the other, we will mean algebraic in a model theoretic sense. We will see later on that, under certain field theoretic assumptions, these two notions agree. From here on, if it is not indicated which definition we mean, the reader may assume that we are in a situation in which the definitions coincide.

Definition III.3.

1. Let $K < F$ be fields and let $a \in K$. Then a is **field-theoretically** algebraic over K, denoted $a \in K^{alg}$, if there exists a non-constant polynomial $P(x) \in K[x]$ such that $P(a) = 0$.

2. Let M be a structure, $A \subset M$, $\bar{a}$ a tuple from M. Then $\bar{a}$ is **model theoretically algebraic** over A in M, denoted $\bar{a} \in \mathrm{acl}_M(A)$, if there is $\varphi(\bar{x}) \in \mathcal{L}_A$ such that $M \models \varphi(\bar{a}) \wedge \exists^{\leq k}\bar{x}\ \varphi(\bar{x})$ for some non-zero $k \in \omega$.

Proposition III.4. *Let K be an algebraically closed field. Let $A \subset K$. Let $\langle A \rangle$ be the subfield of K generated by A. That is, take B to be the ring generated by A, then let $\langle A \rangle$ be $\mathrm{Frac}(B)$, the field of fractions of B. Then we have the following:*

1. *Let $b \in K$. Then $b \in \mathrm{acl}_K(A)$ if and only if $b \in \langle A \rangle^{alg}$.*

2. *There is a unique complete type over A or $\langle A \rangle$ of an element $b \in K$ or in some elementary extension of K such that $b \notin \langle A \rangle^{alg}$.*

Proof. Without loss of generality, we will assume that A is a subring of K.

1. If $b \in \langle A \rangle^{alg}$, then there is a polynomial $P(x) \in K[x]$ such that $P(b) = 0$. By clearing the denominators, $P(x)$ becomes a polynomial with coefficients in A, which can be viewed as a formula $\varphi(x) \in \mathcal{L}_A$. Since $P(x)$ has finitely solutions, $\varphi(x)$ has finitely many solutions or realizations in K, and thus, $b \in \mathrm{acl}_K(A)$.

   For the other direction, first note that $\{P(x) \neq 0 : P(x) \in A[x]\}$ axiomatizes a complete type over A by quantifier elimination. Let $b \notin \langle A \rangle^{alg}$, and consider the complete type of b over A denoted by $\mathrm{tp}_K(b/A)$. Then since $b \notin \langle A \rangle^{alg}$ we have that '$f(x) = 0$' is not in $tp_K(b/A)$ for any $f(x) \in A[x]$. Thus, by the above, $\{P(x) \neq 0 : P(x) \in A[x]\}$ axiomatizes $\mathrm{tp}_K(b/A)$. By compactness, for $\varphi(x) \in p(x)$ such that $\varphi(x)$ is an $\mathcal{L}_A$ formula, there exists $P_1, \ldots, P_r \in A[x]$ such that $K \models \forall x(\bigwedge_{i=1}^{r} P_i(x) \neq 0 \rightarrow \varphi(x))$. But $\bigwedge_{i=1}^{r} P_i(x) \neq 0$ has infinitely many solutions in K, which means $\varphi(x)$ also has infinitely many solutions, and so $b \notin \mathrm{acl}_K(A)$.

2. If $b \in K_1 \succ K$ and $b \notin \langle A \rangle^{alg} < K$ then as we noted in 1., $\mathrm{tp}_K(b/\langle A \rangle)$ is axiomatized by $\{P(x) \neq 0 : P(x) \in \langle A \rangle[x]\}$. We have that $\{P(x) \neq 0 : P(x) \in A[x]\}$ is consistent by compactness.

□

Definition III.5. A field K is called perfect if either $\mathrm{char}(K) = 0$ or, if $\mathrm{char}(K) = p$, then for all $a \in K$ there is $b \in K$ such that $a = b^p$.

Remark III.6. Let K be a field of characteristic p. Then the map $\mathrm{Fr} : K \to K$ such that $\mathrm{Fr}(a) = a^p$ is called the Frobenius map. We know that Fr is an injective field endomorphism since the kernel of Fr is not the whole field (recall that the only ideals of a field F are $\{0\}$ and F). To be perfect means that Fr is also surjective.

Exercise III.7. Let K be a field.

1. If K is finite, then it is perfect.
2. K^{alg} is perfect.
3. Let K be algebraically closed, and let $A \subset K$. Then $\langle A\rangle^{alg}$ is the perfect closure of the field generated by A, where the perfect closure is closed under p^{th} roots.

Fix an algebraically closed field K. We shall naïvely say that a subset $V \subseteq K^n$ is an algebraic variety (sometimes shortened to just variety) if there is a finite set $S \subset K[x_1, \ldots, x_n]$ of polynomials such that for all $\bar{x} \in V$, $f(\bar{x}) = 0$ for all $f \in S$. More specifically, this is the definition of the set of K-rational points of an affine, possibly reducible, variety. An algebraic variety is also called a Zariski closed or an algebraic subset of K^n.

Note that a variety is a special case of a definable set, i.e., a set definable in $(K, +, \cdot, 0, 1)$ possibly with parameters, since we have that $\varphi(\bar{x}) = P_1(\bar{x}) = 0 \wedge \ldots \wedge P_r(\bar{x}) = 0$. Also, by quantifier elimination, any definable subset with parameters in K^n is a finite Boolean combination of varieties.

Proposition III.8. *1. Any ideal $I \subset K[x_1, \ldots, x_n]$ is finitely generated.*

2. Let K be algebraically closed. Then there exists a $1-1$ correspondence between algebraic varieties $V \subset K^n$ and radical ideals I of $K[x_1, \ldots, x_n]$.

3. We have the descending chain condition (DCC) on algebraic varieties: every descending chain of varieties is eventually constant.

Proof. 1. This is also known as the Hilbert Basis Theorem. See any algebra textbook for a proof.

2. We give a $1-1$ correspondence: let $V \subset K^n$ and let $I(V)$ be the ideal $\{f \in K[x] : f(\bar{v}) = 0, \forall \bar{v} \in V\}$. Given an ideal $I \subset K[x]$, let $V(I) = \{\bar{a} \in K^n : f(\bar{a}) = 0 \; \forall f \in I\}$. Since I is finitely generated, I is of the form $(f_1, \ldots, f_r)$. Thus, $V(I) = \{\bar{a} \in K[x] : f_1(\bar{a}) = 0, \ldots, f_n(\bar{a}) = 0\}$, and so $V(I)$ is a variety. Note that $\forall V \subseteq K^n$, $I(V)$ is actually a radical ideal, since for $f^n(\bar{a}) = f(\bar{a})^n = 0 \Rightarrow f(\bar{a}) = 0$. This gives the required $1-1$ correspondence, and so for each $V \subset K^n$ a variety, $V(I(V)) = V$ and for each radical ideal I we have $I(V(I)) = I$.

3. The DCC means that there is no infinite chain of subvarieties of the form $V_1 \supsetneq V_2 \supsetneq \ldots$. Equivalently, the DCC means that if $\{V_j : j \in \beta\}$ for β an ordinal, is a collection of algebraic varieties in K^n, then $\bigcap_j V_j = V_{j_1} \cap \ldots \cap V_{j_r}$ for some $j_1 \ldots, j_r \in \beta$, since the collection $\{\cap_{i<j} V_i : j \in \beta\}$ forms a descending chain. Let $\{V_j : j \in \beta\}$ be a collection of varieties, and let $I_j = I(V_j)$. Then $\bigcup_{j \in \beta} I_j$ is an ideal in $K[\bar{x}]$ so it is finitely generated by some $f_1, \ldots, f_r$. Therefore, we have that $\bigcap_{j \in J} V_j$ is defined by $f_1, \ldots, f_r$. That is, $\bigcap_{j \in J} V_j = \{\bar{a} \in K^n : f_1(\bar{a}) = \ldots = f_r(\bar{a}) = 0\}$. The result follows.

□

Example III.9. The algebraic subvarieties of a field K are K, any finite subset of K, and $\emptyset$. K is a subvariety because it is defined by the vanishing points of the zero polynomial. Similarly, $\emptyset$ is a subvariety defined by any non-zero constant polynomial. Otherwise, if $I(V)$ contains some polynomial $p(x)$ of degree $p \geq 1$, then the solution set is finite and thus, V is finite. On the other hand, let $a_1, \ldots, a_n \in K$. Then it is the solution set for $p(x) = (x - a_1) \ldots (x - a_n)$.

Definition III.10. For $n \in \omega$, we define the *Zariski topology* on K^n to be the topology generated by taking algebraic varieties as closed sets.

Exercise III.11. The Zariski topology on K^n is Noetherian (i.e. satisfies the descending chain condition on closed sets) and compact, but not Hausdorff.

Definition III.12. Let $V \subseteq K^n$. V is called irreducible if we cannot write $V = V_1 \cup V_2$, where V_1, V_2 are non-empty varieties properly contained in V.

Exercise III.13. Let $V \subseteq K^n$ be a variety. Then:

1. V is irreducible if and only if $I(V)$ is prime.
2. K^n is irreducible for all n.

Proposition III.14. *Any variety $V \subseteq K^n$ can be written as a unique non-redundant decomposition $V = V_1 \cup \ldots \cup V_r$ where each V_i is an irreducible variety and where $V_i \not\subset V_j$ (however, $V_i \cap V_j$ may be non-empty) for all i, j.*

Proof. To show existence, if V is irreducible, then we are done. If not, then $V = V_1 \cup V_2$ with V_1, V_2 subvarieties of V. If V_1 or V_2 are reducible we can decompose further. Without loss of generality suppose V_1 is reducible, and this process never stops. Then we have a an infinite descending chain of subvarieties $V \supsetneq W_1 \supsetneq \ldots$, violating the descending chain condition. Thus, the process must stop at some finite level.

To show uniqueness, suppose V admits of two non-redundant decompositions $V = V_1 \cup \ldots V_n$ and $V = W_1 \cup \ldots \cup W_m$. Then $V_i = (V_i \cap W_1) \cup \ldots \cup (V_i \cap W_m)$. Since V_i is irreducible, $V_i = W_j$ for some j. Thus, $n \leq m$. We can repeat the argument showing that $m \leq n$. Thus, $n = m$. Since each $V_i = W_j$ for some j, we have that the two decompositions are identical (modulo order). □

Example III.15. 1. The variety in K^2 defined by $xy = 0$ is reducible with components $x = 0$ and $y = 0$, i.e. the y, x axes, respectively.

2. Suppose $P(x, y)$ is a polynomial in two variables. Then the variety defined by $P(x, y) = 0$ is irreducible if and only if $P(x, y)$ is irreducible as a polynomial.

Proposition III.16. *Let $V \subseteq K^n$ be irreducible. Then $\Sigma(\bar{x}) = \{x \in V\} \cup \{x \notin W : W \subsetneq V \text{ where } W \text{ is a variety}\}$ axiomatizes a complete type $p_V(\bar{x})$ over K, which we call the generic type of V. Note: "$x \in V$" just means $\varphi(x)$ where φ defines V.*

Proof. We need to show two things:

1. $\Sigma_V(\bar{x})$ is consistent with $\mathrm{Th}(K, a)_{a \in K}$ and

2. for any $\varphi(\bar{x})$ over K, either $\Sigma_v(\bar{x}) \models \varphi(\bar{x})$ or $\Sigma_V(\bar{x}) \models \neg\varphi(\bar{x})$ (completeness).

To show 1., note that for all $W_1, \ldots, W_r \subsetneq V$, the set defined by "$x \in V \wedge x \notin W_1 \wedge \ldots \wedge x \notin W_r$" is realized in K. Otherwise, $V = \cup_{i=1}^{r} W_i$, which contradicts the irreducibility of V. Thus, by compactness, $\Sigma_V(\bar{x})$ is consistent.

To show completeness, we know by quantifier elimination that any complete type $q(\bar{x}) \in S_n(K)$ is determined by the open (closed) sets in q. We claim that $\Sigma_V(\bar{x})$ already contains or implies all this information. Let $V_1 \subseteq K^n$ be a variety. Case (i): $V \subseteq V_1$. Then $\Sigma_V(\bar{x}) \vdash \bar{x} \in V_1$. Case (ii): Otherwise $V \cap V_1 \subsetneq V$. Thus, $\Sigma_V(\bar{x}) \vdash \bar{x} \notin V \cap V_1$. Since the formula "$x \in V$" is in $\Sigma_V(\bar{x})$, we have that $\Sigma_V(\bar{x}) \vdash x \notin V_1$, as desired.

□

Remark III.17. The notion of generic type should remind the reader of the notion of *generic point* in algebraic geometry. For an algebraic variety V, the scheme associated to V is the collection of prime ideals of $K[V]$, denoted by $Spec(K[V])$ where $K[V] = K[\bar{x}]/I(V)$ is called the coordinate ring of V. Here, the maximal ideals of $K[V]$ are called the closed points of $Spec(K[V])$ and correspond precisely to the points of V. We have a correspondence between the prime ideals of $K[V]$, irreducible subvarities of V, and the types $q(\bar{x}) \in S_n(K)$ extending $\bar{x} \in V$. Thus, what is called the generic point of $Spec(K[\bar{x}])$ is what we previously called $p_V(\bar{x})$.

Definition III.18. Let K be an algebraically closed field. Let F be subfield of K. Let $\{a_i : i \in I\} \subseteq K$. We call $\{a_i : i \in I\}$ algebraically independent over F if $a_i \notin \mathrm{acl}(F \cup \{a_j : j \neq i\})$ for all $i \in I$.

Fact III.19. *1. Suppose $\{a_\alpha : \alpha < \kappa\} \subseteq K$ and $\alpha_\beta \notin \mathrm{acl}(F \cup \{a_\alpha : \alpha < \beta\})$ for all $\beta < \kappa$. Then $\{a_\alpha : \alpha < \kappa\}$ is algebraically independent over F. This essentially follows from strong minimality.*

2. *Let $\bar{a} = (a_\alpha : \alpha < \lambda) \subseteq K$. By a basis of $\bar{a}$ over F, we mean a maximal algebraically independent subtuple of $\bar{a}$ over F. If $\overline{a_0}$ is such a basis, then $\overline{a} \subseteq \operatorname{acl}(K, \overline{a_0})$. Any two bases of $\overline{a}$ over F have the same cardinality. This cardinality is called the transcendence degree, denoted by $tr.deg(\overline{a}/F)$ or $tr.deg(F(\overline{a})/F)$.*

Definition III.20. Let K be an algebraically closed field.

1. $V \subseteq K^n$ an irreducible variety. Let $\bar{a}$ realize $p_V(\bar{x})$ in $K_1 \succ K$. Then $\dim(V)$, the dimension of V, is defined to be $tr.deg(\bar{a}/F)$.

2. If $V \subseteq K^n$ is an arbitrary variety, then $\dim(V)$ is defined as the maximum of the dimensions of the irreducible subvarieties of V.

Here, we recall the definition of *Morley Rank*. Let M be a saturated model. Let X be a definable set in M. The Morley Rank of X, denoted $\mathrm{RM}(X)$ is a way of measuring the complexity of X. Morley rank is ordinal valued and is meant to generalize the notion of dimension. Indeed, in the case of varieties, the two notions agree .We define $\mathrm{RM}(X)$ inductively:

1. $\mathrm{RM}(X) \geq 0$ if $X \neq \emptyset$.

2. $\mathrm{RM}(X) \geq \alpha + 1$ if there exists definable $X_i \subseteq X$, for $i = 1, 2, 3, \ldots$ such that $X_i \cap X_j = \emptyset$ for all $i \neq j$ and $\mathrm{RM}(X_i) \geq \alpha$.

3. $\mathrm{RM}(X) = \alpha$ if $\mathrm{RM}(X) \geq \alpha$ but $\not\geq \alpha + 1$.

4. $\mathrm{RM}(X) \geq \delta$ if $\mathrm{RM}(X) \geq \alpha$ for all $\alpha < \delta$.

5. If $\mathrm{RM}(X) > \alpha$ for every α, we say $\mathrm{RM}(X) = \infty$.

Exercise III.21. In DLO, $\mathrm{RM}(x = x)$ is ∞. In ACF, $\mathrm{RM}(x = x) = 1$.

Fact III.22. *1. Let $V \subseteq K^n$ be an irreducible variety. Then $\dim(V)$ equals "Krull dimension" of V, i.e. the maximum length n of a chain $\emptyset \neq V_1 \subsetneq \ldots \subsetneq V_n = V$ of irreducible varieties. Note that included in this is the following: Suppose $V \subseteq K^n$ is irreducible. Let $K_1 \succ K$, $\bar{a_1} \in K_1^n, \bar{a_1} \in V(K_1)$ such that $\dim(V) = tr.deg(\bar{a_1}/K)$. Then $\bar{a_1}$ realizes $p_v(\bar{x})$.*

2. For a variety $V \subseteq K^n$, $\dim(V) = RM(V)$.

Finally, we address definability issues. Let K be algebraically closed and F a subfield of K. Let $V \subseteq K$ be a variety. There are at least three distinct ways in which V is said to be defined over K:

1. V is defined by some formula $\varphi(\bar{x})$ with parameters from F in K;

2. V is defined by some $P_1 = P_2 = \ldots = P_r = 0$ where each $P_i(\bar{x}) \in F[\bar{x}]$;

3. $I(V)$ is generated by polynomials over F. That is, there exists an ideal $I \subset K[\bar{x}]$ such that $I(V) = I \otimes_{F[\bar{x}]} K[\bar{x}]$.

Fact III.23. *When F is perfect, 1., 2., and 3. coincide.*

Proof. Exercise. $\square$

Note that in general, the three do not coincide if F is not perfect. Consider the following counterexample: Let $\mathrm{char}(K) = p$, and let F be a subfield of K. Let $a \in K \setminus F$ and $b = a^p \in F$. Then $\{a\}$ is defined over F in sense of 1. and 2. by the formula $x^p = b$, but it is not in the sense of 3..

Let K be algebraically closed. We can consider the notion of F-irreducibility for a variety V defined over F for F a perfect subfield of K. That is, V (over F) is F-irreducible if we cannot write V properly as $V_1 \cup V_2$ where V_1, V_2 are also defined over F. We have that any variety over K is uniquely $V = V_1 \cup \ldots \cup V_r$ with V_i F-irreducible. Ultimately, F-irreducible varieties over F correspond to complete types over F.

Proposition III.24. *Let $F < K$ be algebraically closed and V a variety over F. Then V is irreducible if and only if V is F-irreducible.*

Proof. We know immediately that V irreducible implies V is F-irreducible. Thus, suppose V were not irreducible. Then $V = V_1 \cup V_2$ with V_1, V_2 defined by equations over K. We can find the right coefficients in F such that V_1 and V_2 are also defined over F since V is a variety over F and $F < K$, showing that V is F-reducible.

$\square$

This motivates the followng definition:

Definition III.25. A variety V defined over a field F is said to be *absolutely irreducible* if it is irreducible over $K > F$ for K algebraically closed.

Chapter IV

Pseudofinite Fields

In this chapter, we summarize some of the model theory of pseudofinite fields. In particular, we study the theory T_f of finite fields, including the structure of infinite models, axiomatization, possibility or impossibility of quantifier elimination, and (non-standard) cardinalities of definable sets.

Our first goal will be to isolate the first-ordered properties shared by all finite fields. In particular we will see that infinite pseudofinite fields are precisely those fields K such that

1. K is perfect,
2. For every $n \in \omega$, there is a unique degree n extension of K (in some fixed algebraic closure of K),
3. K is pseudo-algebraically closed: every variety V defined over K has a K-rational point.

and, in fact, these properties are expressable in a first-order way. Note that property 2. is equivalent to asserting that the absolute Galois group $\mathrm{Gal}(K^{alg}/K)$ is the profinite completion of the integers, $\hat{\mathbb{Z}}$. Recall that $\hat{\mathbb{Z}}$ is defined to be $\varprojlim \mathbb{Z}/n\mathbb{Z}$, where the connecting maps are just the natural " $\bmod m$" maps $\mathbb{Z}/n\mathbb{Z} \to \mathbb{Z}/m\mathbb{Z}$ when m divides n.

Exercise IV.1. Consider $(a_n)_{n\in\omega} \in \prod_{n\in\omega} \mathbb{Z}/n\mathbb{Z}$. Then $(a_n)_{n\in\omega} \in \hat{\mathbb{Z}}$ if and only if whenever m divides n, $a_m \equiv a_n \mod n$.

Perfect fields satisfying property 2. are sometimes called *quasi-finite fields.* The requirement that K above be infinite is just because finite fields, while not pseudo-algebraically closed, are so in an asymptotic sense, which we will see. We begin with some background information.

Proposition IV.2. *1. Any finite field of characteristic $p > 0$ has cardinality p^n for some $n \geq 1$.*

2. For each $q = p^n$, the multiplicative group of $\mathbb{F}_q$ is cyclic.

3. *For each $n \in \omega$, there is a unique finite field of cardinality $q = p^n$ which we shall denote by $\mathbb{F}_q$.*

Proof. 1. Let F be a finite field of characteristic p. Then $\mathbb{F}_p \leq F$ and so F is a finite dimensional $\mathbb{F}_p$-vector space, say of dimension n. Therefore $|F| = p^n$.

2. Let G_q be the multiplicative subgroup of $\mathbb{F}_q$. G_q is a commutative group of order $q-1$. For every $a \in G_q$, there is $m_a \in \omega$ such that $a^{m_a} = 1$. Let $m = \mathrm{lcm}\{m_a : a \in G_q\}$. Then for every $a \in G_q$, $a^m = 1$. Since $\mathbb{F}_q$ is a field, there are at most m solutions to $x^m = 1$ in $\mathbb{F}_q$ and so $m \geq (q-1)$ (since clearly 0 is not a solution). By Lagrange's theorem, $m \mid (q-1)$ and hence $m = (q-1)$. Thus, every element of G_q satisfies $x^{q-1} = 1$. Since $q = p^n$ is a power of a prime, G_q is cyclic by the classification theorem for finite commutative groups.

3. Let K be some algebraically closed field containing $\mathbb{F}_p$ (e.g. the algebraic closure of $\mathbb{F}_p$) and let $q = p^n$ for some $n \geq 1$. Let

$$P(x) = x^q - x = x(x^{q-1} - 1).$$

Recall that, since the formal defivative $P'(x) = -1$ is non-zero, $P(x) = 0$ has no multiple roots. Therefore, it has exactly q distinct roots in K. Let S be the set of roots of $P(x) = 0$. It is an easy exercise to show that S is closed under multiplication, addition, and inverses, so S is a field of characteristic cardinality q. To show uniqueness, suppose F is some other field of cardinality q and so $\mathbb{F}_p < F < K$. Note that every $a \in F$ satisfies $x^q = x$ as in part 2.. Therefore $x^q - x$ vanishes on F, so F equals the set of solutions of $x^q - x$ in K.

$\square$

Corollary IV.3. *If F is a finite field, then, for all n, F has exactly one extension L of degree n. In fact, given an algebraically closed $K > F$, there exists unique $F < L < K$ such that $[L : F] = n$ (if F is of cardinality p^m then L has to be the unique subfield of K of cardinality p^{mn}). In particular, $\mathbb{F}_{p^n}$ is a subfield of $\mathbb{F}_{p^m}$ if and only if n divides m.*

The algebraic closure of any F is the unique (up to isomorphism over F) field $L \geq F$ such that L is algebraically closed and L is algebraic over F, i.e. for all $a \in L$, $P(a) = 0$ for some polynomial P over F. If K is a given algebraically closed field containing F, then

$$\{a \in K : a \text{ algebraic over } F\}$$

is an algebraically closed field and is the algebraic closure of F.

Corollary IV.4. *Fix $\mathbb{F}_p \leq K$, K algebraically closed.*

1. *The algebraic closure of $\mathbb{F}_p$ inside K is $\mathbb{F}_p^{alg} = \bigcup_{n\in\omega} \mathbb{F}_{p^n}$.*

2. *The absolute Galois group of* $\mathbb{F}_p$*, denoted* $\mathrm{Gal}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$*, is* $\hat{\mathbb{Z}}$*.*

Proof. 1. Clear from above. Note that this is not the union of an ascending chain of fields. However, it is clear that we may also write $\mathbb{F}_p^{alg} = \bigcup_{n\in\omega} \mathbb{F}_{p^{n!}}$, which is the union of an ascending chain of fields.

2. Note that for every n, the extension $\mathbb{F}_{p^n}/\mathbb{F}_p$ is Galois of degree n since finite fields are perfect. By definition, for every m, the field $\mathbb{F}_{p^m}$ is precisely the field fixed by Fr^m, the m-th power of the Frobenius map $\mathrm{Fr}(x) = x^p$, and so in particular, $\mathrm{Fr}^m \in \mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p)$. Now, Fr^m fixes $\mathbb{F}_{p^n}$ if and only if n divides m. Therefore Fr has order n in $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p)$. Since $\mathbb{F}_{p^n}/\mathbb{F}_p$ is Galois of degree n, $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p) = \langle \mathrm{Fr} \rangle \cong \mathbb{Z}/n\mathbb{Z}$, the cyclic group with elements $\{\mathrm{id}, \mathrm{Fr}, \mathrm{Fr}^2, \ldots, \mathrm{Fr}^{n-1}\}$.

   Let $\sigma \in \mathrm{Gal}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$. Since $\mathbb{F}_p^{alg}$ is the union over all n of $\mathbb{F}_{p^n}$, σ is determined by the restriction to each $\mathbb{F}_{p^n}$. Since $\mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p) \cong \mathbb{Z}/n\mathbb{Z}$, we have a natural map from $\mathrm{Gal}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$ to $\hat{\mathbb{Z}}$ defined by $\sigma \mapsto (a_n)_{n\in\omega}$ where $\sigma \restriction_{\mathbb{F}_{p^n}} = \mathrm{Fr}^{a_n}$ where $a_n \in \mathbb{Z}/n\mathbb{Z}$. We leave it as an exercise to show that this map is in fact a group homomorphism. It is clear that this map is injective, since $\sigma \mapsto (0, 0, \ldots)$ if and only if $\sigma \restriction_{\mathbb{F}_{p^n}} = \mathrm{Fr}^n$ for every $n \in \omega$. To see that this map is surjective, note that if $(a_n)_{n\in\omega} \in \hat{\mathbb{Z}}$, then the sequence of automorphisms $\mathrm{Fr}^{a_n} \in \mathrm{Gal}(\mathbb{F}_{p^n}/\mathbb{F}_p)$ determines an automorphsims $\sigma \in \mathrm{Gal}(\mathbb{F}_p^{alg}/\mathbb{F}_p)$. This is left as an exercise.

□

Observe that, an easy adaptation of the arguement of part 2. from above, for any finite field F, $\mathrm{Gal}(F^{alg}/F) = \hat{\mathbb{Z}}$. As we will see, this property, equivalent to F having exactly one extension of degree n for all n, is expressable in a first-order way. This suggests that this property should be expressed in our theory T_f of finite fields.

Proposition IV.5. *There is an axiom schema* $(\varphi_n)_{n\in\omega}$ *in the language of unitary rings such that a field* $F \models \varphi_n$ *if and only if* F *has a unique algebraic extension of degree* n*.*

Proof. We follow the proof given in [1] Let $n \in \omega$. For every $1 \le m < n$ consider the formula $\psi_m(\bar{y})$ in variables $y_0, \ldots, y_{n-1}$ which expresses the fact that $(\forall z_0, \ldots, z_{m-1}, w_0, \ldots, w_{n-m-1})$

$$\left(X^n + \sum_{i=0}^{n-1} y_i X^i\right) \neq \left(X^m + \sum_{j=0}^{m-1} z_j X^j\right)\left(X^{n-m} + \sum_{h=0}^{n-m-1} w_h X^h\right)$$

i.e. that the polynomial $X^n + \sum_{i=0}^{n-1} y_i X^i$ is not expressible as the product of a polynomial of degree m with a polynomial of degree $n - m$. Let ψ_n be the the sentence $\exists \bar{y} \bigwedge_{m=1}^{n-1} \psi_m(\bar{y})$. This sentence then asserts that there is a irreducible polynomial of degree n. That is $F \models \psi_n$ if and only if F has an algebraic extension of degree n.

To finish the proof we need to be able to express the uniqueness of the field extension of degree n. Recall that if F is a field and $P(x) = \sum_i c_i x^i$ is an irreducible polynomial of degree n with tuple of coefficients $\bar{c} = (c_0, \ldots, c_{n-1}, 1)$ in F, then for $\alpha \in F^{alg}$ a solution of $P(x) = 0$, the field extension $F(\alpha)$ of degree n is an F-vectorspace with basis $\{1, \alpha, \ldots, \alpha^{n-1}\}$. Note that $F(\alpha)$ is the unique estension of degree n if and only if, for every irreducible polynomial $Q(x)$ over F of degreen n, $Q(x) = 0$ has a solution in $F(\alpha)$. Thus, the strategy is to show that the field $F(\alpha)$ is uniformly interpretable in F using only the parameters $\bar{c}$ and then to express that for any field extension of degree n must be the same as $F(\alpha)$.

We interpret $F(\alpha)$ as the F-vector space F^n with basis $\{1, \alpha, \ldots, \alpha^{n-1}\}$. The field F is interpreted as the set $\{(a, 0, 0, \ldots, 0) : a \in F\}$, α is represented as $(0, 1, 0, \ldots, 0)$ and so on. Now, observe that multiplication by α induces the linear map

$$L_\alpha = \begin{pmatrix} 0 & 0 & \ldots & 0 & -c_0 \\ 1 & 0 & \ldots & 0 & -c_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & -c_{n-1} \end{pmatrix}$$

since $\alpha^n = \sum_{i=0}^{n-1} -c_i \alpha^i$. Thus, multiplication by α^k is given by the linear map L_α^k. We can therefore interpret multiplication in $F(\alpha)$ as

$$(a_0, \ldots, a_{n-1}) \times (b_0, \ldots, b_{n-1}) \equiv \left(a_0 I + a_1 L_\alpha + \ldots + a_{n-1} L_\alpha^{n-1}\right) \begin{pmatrix} b_0 \\ b_1 \\ \vdots \\ b_{n-1} \end{pmatrix}.$$

Observe that this multiplication is definable using just $\bar{c}$ as a parameter. Let $\theta_n(\bar{y})$ be for fomula $\bigwedge_{m=1}^{n-1} \psi_m(\bar{y})$ from before expressing irreducibility. Clearly $F \models \theta_n(\bar{c})$. Suppose that $F \models \theta_n(\bar{b})$ also. That is, suppose that $Q(x) = \sum_{i=0}^{n} b_i x^i$ is irreducible over F. Then the field extension generated by a root of $Q(x) = 0$ is the same as $F(\alpha)$ if and only if there is $\beta \in F(\alpha)$ with $Q(\beta) = 0$. That is, if and only if there is $(a_0, \ldots, a_{n-1}) \in F^n$ such that

$$\sum_{i=0}^{n} b_i (a_0, \ldots, a_{n-1})^i = (0, \ldots, 0)$$

where the power $(a_0, \ldots, a_{n-1})^i$ is defined as above. Using this, let $\chi(\bar{b}, \bar{c})$ be the formula expressing that if $\theta_n(\bar{b})$ and $\theta_n(\bar{c})$, then the irreducible polynomials with coefficients $\bar{b}$ and $\bar{c}$ each generate the same degree n extensions of F. Then φ_n is $\psi_n \wedge \forall \bar{x}, \bar{y} \chi(\bar{x}, \bar{y})$. $\square$

What more can we say about T_f?

Fact IV.6 (Lang-Weil Estimates). *Fix n, d. Then there exists a constant $C(n,d) = C \in \mathbb{N}$ such that for any finite field $\mathbb{F}_q$ and (absolutely) irreducible variety V $(\subseteq (\mathbb{F}_q^{alg})^n)$ over $\mathbb{F}_q$ defined by finitely many polynomials in $\mathbb{F}_q[x_1, \ldots, x_n]$ of degree $\leq d$,*

$$||V(\mathbb{F}_q)| - q^{\dim(V)}| \leq C q^{\dim(V) - \frac{1}{2}}.$$

Remark IV.7. 1. $V(\mathbb{F}_q) = \{(x_1, \ldots x_n) \in \mathbb{F}_q^n : (x_1, \ldots x_n) \in V\}$

2. So if $\sqrt{q} > C$, then $V(\mathbb{F}_q) \neq \emptyset$ since

$$\left| \frac{|V(\mathbb{F}_q)|}{q^{\dim(V)}} - 1 \right| \leq \frac{C}{\sqrt{q}}.$$

The goal here is to have a suitable translation of the Lang-Weil estimates for a (saturated) pseudofinite field of cardinality $q \in \mathbb{N}^*$ (and so $q \geq n$ for all $n \in \mathbb{N}$). Indeed, we have the following:

Fact IV.8. *The pseudofinite analogue of the Lang-Weil Estimates is true by transfer: For V an absolutely irreducible variety over F defined by polynomials in $F[x_1, \ldots x_n]$ of degree $\leq d$,*

$$\Big| |V(F)| - q^{\dim(V)} \Big| \leq C q^{\dim(V) - \frac{1}{2}}$$

for some $C = C(n,d) \in \mathbb{N}$.
Indeed, $\left| \frac{|V(F)|}{q \dim(V)} - 1 \right| \leq \frac{C}{\sqrt{q}}$ is infinitesimal, so

$$\mu_{\dim(V)}(V) := \mathrm{st}\left(\frac{|V(F)|}{q^{\dim(V)}} \right) = 1.$$

Exercise IV.9. If $V = \mathbb{A}$ then $\dim(V) = n$ and $|V(F)| = |\mathbb{F}^n|$.

Remark IV.10. Consider $x^2 + y^2 - 1 = 0$. In a field F, we may think number of solutions is approximately $\geq 2|F|$. But we do not expect the same for $y - x^2 = 0$, because not every y is a square.

The following is a crucial fact:

Fact IV.11. *We assume that we are working in the theory ACF. Fix $n, d, r \in \mathbb{N}$. Let $P_1(\bar{x}), \ldots, P_r(\bar{x})$ be polynomials of degree $\leq d$ in variables $\bar{x} = (x_1, \ldots, x_n)$ with $P_i(\bar{x})$ having tuple of coefficients $\bar{a}_i$ all in some algebraically closed field K. Let V be the variety defined by $P_1 = P_2 = \ldots = P_r = 0$.*

1. *Let $D \leq n$. Then there is a formula $\theta_D(\bar{y}_1, \ldots, \bar{y}_n)$ such that*

$$\dim(V) = D \text{ iff } K \models \theta_D(\bar{a}_1, \ldots, \bar{a}_n)$$

2. *There is $\theta_2(\bar{y}_1, \ldots, \bar{y}_n)$ such that*

$$V \text{ is (absolutely) irreducible iff } \theta_2(\bar{a}_1, \ldots, \bar{a}_n)$$

Recall that T_f = theory of finite fields in the language of unitary rings. We will let T_f^* be the theory in the same language axiomatised by the axioms for fields and:

1. F is perfect,
2. F has exactly one extension of degree n for all n,
3. Suppose V is an absolutely irreducible variety over F defined by r polynomials in $x_1, \dots, x_n$ of degree $\leq d$ and if $|F| > C(n,d)^2$ then $V(F) \neq \emptyset$. By IV.11 this is expressible.

By earlier lemmas, $T_f \models T_f^*$, i.e. every finite field is a model of T_f^*. Our goal now is to show that $T_f^* \models T_f$, i.e. $T_f^* = T_f$.

Remark IV.12. Note that the theory of the class of infinite models of T_f^* is axiomatized by 1. from above, and the axiom scheme expressing that any (absolutely) irreducible variety V defined over F contains an F-rational point (F is pseudo-algebraically closed).

Exercise IV.13. If F is a finite field, then some absolutely irreducible variety over F has no F-point.

Therefore, we wish to show that $T_f^\infty = (T_f^*)^\infty$.

We now give a summary of model theory of pseudofinite fields. For more details see [3], [2].

Lemma IV.14. *Let K, L be models of $(T_f^*)^\infty$. Suppose K, L have a common subfield k which is relatively algebraically closed in K and L (in the field theoretic sense). That is, whenever $b \in K$ is algebraic over k (in the field theoretic sense) then $b \in k$ and likewise for L, i.e. $k^{alg} \cap K = k$, $k^{alg} \cap L = k$. (Note that $k \prec K$ implies that k relatively algebraic in K.) Then $K \equiv_k L$, i.e.*

$$\mathrm{Th}(K,a)_{a\in k} \equiv \mathrm{Th}(L,a)_{a\in k}.$$

Proof. We may assume K, L are $(|k|+\omega)^+$-saturated. Now we try to show that the collection of isomorphisms f between small relatively algebraically closed subsets E_1 of K, E_2 of L, $f : E_1 \cong E_2$ such that $f \restriction k = \mathrm{id}$, has the back and forth property.
Let $a \in K$. We want to extend $\mathrm{id} : k \cong k$ to some relatively algebraically closed E containing $k \cup \{a\}$. Let E be the relative algebraic closure of $k \cup \{a\}$ in K, $E = k(a)^{alg} \cap K$. Want to find $f : E \simeq_k f(E)$ relatively algebraically closed in L. Let $\bar{b} \subset E$. Let $I(\bar{b}/k)$ be the set of polynomials $P(x_1, \dots, x_n)$ over k such that $P(\bar{b}) = 0$. Let $V = V(I(\bar{b}/k))$. As k is relatively algebraic in K, V is (absolutely) irreducible. So $V(L) \neq \emptyset$ by since L is pseudo-algebraically closed. For the same reason, V has a generic point (over k) in L, say $\bar{c}$. i.e.

$$I(\bar{c}/k) = I(\bar{b}/k)$$
$$\mathrm{qftp}(\bar{c}/k) = \mathrm{qftp}(\bar{b}/k)$$

So $f : k(\bar{b}) \to k(\bar{c})$, $f(\bar{b}) = \bar{c}$, $f \restriction k = \mathrm{id}$ is an isomorphism. i.e. for all $\bar{b} \in E$, $\mathrm{qftp}(\bar{b}/k)$ can be realized in L.
Let $\vec{b}$ enumerate all of E. By saturation of L, one can realize $\mathrm{qftp}(\vec{b}/k)$ in L. This yields an embedding $f : E \hookrightarrow L$ over k. It is left as an exercise to show that we may choose f such that $f(E)$ is relatively algebraically closed in L using the fact that any model of $(T_f^*)^\infty$ has absolute Galois group $\hat{\mathbb{Z}}$. □

Corollary IV.15. *Suppose $K \subset L$ are models of $(T_f^*)^\infty$ and $K^{alg} \cap L = K$. Then $K \prec L$. (Take $k = K$ as in IV.14.)*

Definition IV.16. Let K be any field. By $\mathrm{Abs}(K)$, the field of "absolute numbers" of K, we mean $k_0^{alg} \cap K$, where $k_0 < K$ is the prime field $\mathbb{F}_q$ or $\mathbb{Q}$.

Corollary IV.17. *If $K, L \models (T_f^*)^\infty$, then $K \equiv L$ iff $\mathrm{Abs}(K) \cong \mathrm{Abs}(L)$.*

Proof. Exercise. □

Question What are the possibilities for $\mathrm{Abs}(K)$, $K \models (T_f^*)^\infty$?

Remark IV.18.

1. Suppose $K \models (T_f^*)^\infty$ and E is relatively algebraically closed in K. Then E has at most one extension of degree n, for all n, because if E_1, E_2 were different such extensions, then KE_1, KE_2 would be different extensions of K of degree n.

2. $\mathrm{Abs}(K)$ has at most one extension of degree n for all n.

Lemma IV.19. *Let $k_0 = \mathbb{F}_p$ or $\mathbb{Q}$, and $E \subseteq k_0^{alg}$ such that E has at most one extension of degree n for all n. Then there exists a non-principal ultraproduct of finite fields K such that $\mathrm{Abs}(K) \cong E$.*

As an example, consider the special case of $E = \mathbb{F}_p$. Let P be set of all primes and let $\mathcal{U}$ be a non-principal ultrafilter on P. Let $K = \prod_{r \in P} \mathbb{F}_{p^r}/\mathcal{U}$. Then $\mathrm{Abs}(K) = \mathbb{F}_p$.

Exercise IV.20. Prove Lemma IV.19. (HARD!)

Corollary IV.21. $T_f = T_f^*$.

Proof. We observed earlier that $T_f \models T_f^*$ and so we want to show that $T_f^* \models T_f$. If $K \models T_f^*$ and K is finite, then $K \models T_f$. So we may assume K is infinite, i.e. $K \models (T_f)^\infty$. Let $E = \mathrm{Abs}(K)$. Then E has at most one extension of degree n, for all n. By IV.19, there exists L, a non-principal ultraproduct of finite fields, such that $\mathrm{Abs}(L) = E$. Therefore, $L \models (T_f^*)^\infty$ and $E \subseteq K, L$, with E relatively algebraically closed in K and L. By IV.17 (or IV.14), $K \equiv L$, so $K \models T_f$. □

Therefore, we have show that $T_f = T_f^*$ and thus a pseudofinite field is precisely a model of T_f. In particular, a strictly pseudofinite field is a model of T_f^∞.

The theory T_f^∞ is incomplete. However, by Lemma IV.19, we know the completions of T_f^∞. We note that T_f does not eliminate quantifiers. For example, the formula $\varphi(x)$ given by "$(\exists y)x = y^2$" is not equivalent to a quantifier-free formula modulo T_f. If it were, the for a model $M \models T_f^\infty$, $\varphi(M)$ would be finite or cofinite. However, for a finite field, precisely half of the non-zero elements are squares and so, in a strictly pseudofinite field, $\varphi(x)$ defines an infinite, coinfinite set.

Though we do not get full quantifier elimination, we can expand our language in such a way that we can eliminate quantifiers up to existential formulas. We expand the language by new constant symbols

$$c_{i,n}, \quad n \geq 2, \quad 0 \leq i \leq n-1$$

and get the theory $T_{f,\bar{c}}$ by adding to T_f the axioms expressing that

$$x^n + c_{n-1,n}x^{n-1} + \ldots + c_{1,n}x + c_{0,n}$$

is irreducible for every n. Any pseudofinite field F has an expansion to model of $T_{f,\bar{c}}$, by the fact that every psdeudofinite field has a unique algebraic extension of degree n.

Lemma IV.22. *Let $K, L \models T_{f,\bar{c}}^\infty$, and let k be a common substructure of K and L (so containing interpretations of the $c_{i,j}$). If k is relatively algebraically closed in K then k is also relatively algebraically closed in L.*

Proof. Suppose k is relatively algebraically closed in K but not relatively algebraically closed in L. Since k is not relatively algebraically closed in L, there exists a finite extension $k < k_1 < L$, of degree n for some $n \in \omega$ over k. By an earlier observation k_1 is the unique extension of degree n. Since k is relatively algebraically closed in K, the compositum Kk_1 is an extension of K of degree n and so generated by a solution of

$$P(x) = x^n + c_{n-1,n}x^{n-1} + \ldots + c_{1,n}x + c_{0,n} = 0.$$

Since Kk_1 is generated by the same element which generated k_1 over k, there is a solution of $P(x) = 0$ in $k_1 < L$, which contradicts the fact that $P(x)$ is irreducible.

□

Recall that a theory T has quantifier elimination if and only if for every $M \models T$ and finite tuple $\bar{a} \subset M$, $T \cup \mathrm{qftp}(\bar{a})$ is complete (in the language expanded by constants for $\bar{a}$). Using a similar idea, we show that we can eliminate quantifiers up to existential formulas modulo the theory.

Corollary IV.23. *Let $K \models T_{f,\bar{c}}^\infty$. Let $k < K$ be relatively algebraically closed. Let $\vec{a}$ enumerate k. Then $T_{f,\bar{c}}^\infty \cup \mathrm{qftp}(\vec{a})$ is complete.*

Proof. Let $(L, \vec{a}) \models T_{f,\bar{c}}^\infty \cup \mathrm{qftp}(\vec{a})$. By IV.22, k is relatively algebraically closed in L too. So by IV.14, $K \equiv_k L$, i.e. $\mathrm{tp}_L(\vec{a}) = \mathrm{tp}_K(\vec{a})$. □

Compactness yields the following:

Corollary IV.24. *For every formula $\varphi(\bar{x}) \in L(T^{\infty}_{f,\bar{c}})$ there is a positive, quantifier-free $\psi(\bar{x}, \bar{y})$ such that*

$$T^{\infty}_{f,\bar{c}} \models \forall \bar{x}(\varphi(\bar{x}) \leftrightarrow \exists \bar{y} \psi(\bar{x}, \bar{y}))$$

and therefore $T^{\infty}_{f,\bar{c}}$ is model complete. Furthermore, for some $N \in \mathbb{N}$,

$$T^{\infty}_{f,\bar{c}} \models \forall \bar{x} \exists^{\leq N} \bar{y} \psi(\bar{x}, \bar{y}). \tag{IV.1}$$

Proof. Let $K \models T^{\infty}_{f,\bar{c}}$. For any tuple $\bar{b} \in K$, $\mathrm{acl}(\bar{b}) \subseteq K$ is a relatively algebraically closed subfield of K and so $T^{\infty}_{f,\bar{c}} \cup \mathrm{qftp}(\mathrm{acl}(\bar{b}))$ is complete. Let $\varphi(\bar{x})$ be a formula. Then, for any $\bar{b} \in K$, $K \models \varphi(b)$ if and only if there is quantifier free $\psi(\bar{x}, \bar{y})$ and $\bar{d} \in \mathrm{acl}(\bar{b}) \subset K$ such that $K \models \psi(\bar{b}, \bar{d})$. Therefore,

$$T^{\infty}_{f,\bar{c}} \models \forall \bar{x}(\varphi(\bar{x}) \leftrightarrow \exists \bar{y} \psi(\bar{x}, \bar{y})).$$

Here, we may assume that $\psi(\bar{x}, \bar{y})$ is positive, since for any occurence of a subformula of the form $P(\bar{z}) \neq 0$ for P a polynomial, we may replace it with $\exists r(rP(\bar{z}) = 1)$. The bound IV.1 follows from the algebraicity of $\bar{d}$ over $\bar{b}$. $\square$

Essentially, IV.24 says that every definable set is the image of a variety under some N-to-1 map. This result gives us some idea of what definable sets in *finite* fields look like. Roughly, suppose that $F \leq F'$ are finite fields, $\bar{a} \in F$ and $\varphi(\bar{x}, \bar{a})$ is a formula with parameters $\bar{a} \in F$. What can we say about $|\varphi(\bar{x}, \bar{a})(F')|$? By IV.24, we have such a positive, quantifier-free formula $\psi(\bar{x}, \bar{a}, \bar{z})$. Let us assume that $\psi(\bar{x}, \bar{a}, \bar{z})$ defines (in some larger algebraically closed field) an absolutely irreducible variety over F and that in (IV.1) we have "$= N$" rather than "$\leq N$". Then, for a large enough finite field $F' > F$ we get

$$|\varphi(\bar{x}, \bar{a})(F')| = \frac{1}{N} |\psi(\bar{x}, \bar{a}, \bar{z})(F')|$$

using the fact that $\varphi(\bar{x}, \bar{a})$ defines a set which is the image of a variety under an N-to-1 map. By the Lang-Weil Estimates (IV.6), $|\psi(\bar{x}, \bar{a}, \bar{z})|$ is within $C |F'|^{d-1/2}$ of $|F'|^d$ where d is the dimension of the variety defined by $\psi(\bar{x}, \bar{a}, \bar{z})$, and so we get a good description of $|\varphi(\bar{x}, \bar{a})(F')|$ as well. The upshot (with some additional work) is the following:

Theorem IV.25 (Chatzidakis, van den Dries, MacIntyre [3])**.** *Let $\varphi(\bar{x}, \bar{y})$ be a formula in the language of rings with $\bar{x} = (x_1, \ldots, x_m)$ and $\bar{y} = (y_1, \ldots, y_n)$. Then there is a positive constant C and a finite set D of pairs (d, μ) with $0 \leq d \leq n$ and $\mu \in \mathbb{Q}^{>0}$ or $(d, \mu) = (0, 0)$ and for every $(d, \mu) \in D$, there is an L_{ring}-formula $\varphi_{d,\mu}(\bar{y})$ such that for every finite field $\mathbb{F}_q$ and every $\bar{a} \in \mathbb{F}_q^n$, exactly one of the $\varphi_{d,\mu}(\bar{y})$ holds of $\bar{a}$, and, if $\mathbb{F}_q \models \varphi_{d,\mu}(\bar{a})$, then*

$$\left| |\varphi_{d,\mu}(\bar{x}, \bar{a})(\mathbb{F}_q)| - \mu q^d \right| < C q^{d-1/2}.$$

This theorem has a natural interpretation in the non-standard setting. Fix $\varphi(\bar{x}, \bar{y})$ a formula in the language of rings and let F be a pseudofinite fields of non-standard cardinality $q \in \mathbb{N}^*$. Then, for every $\bar{a} \in F^n$, there is a unique $(d, \mu) \in D_\varphi$ such that $\varphi_{d,\mu}(\bar{a})$ holds. Moreover, if $F \models \varphi_{d,\mu}(\bar{a})$, then

$$\operatorname{st}\left(\frac{|\varphi(\bar{x}, \bar{a})(F)|}{q^d}\right) = \mu.$$

That is, to an infinite definable set $\varphi(\bar{x}, \bar{a})(F)$ in an pseudofinite field F, we are associating a dimension d, and a measure, μ.

Exercise IV.26. T_f^∞ is a simple theory and, for $\varphi(\bar{x}, \bar{a})$ as above, d is the SU-rank of $\varphi(\bar{x}, \bar{a})$.

In [12], Tao used IV.25 to prove a regularity lemma for definable graphs in pseudofinite fields.

Open Question: Real closed fields, pseudofinite fields, p-adically closed fields are examples of "nice fields" in the sense of having *almost* quantifier elimination: every formula is equivalent to a formula with finite fibres (existential, plus some other stuff?) as in IV.24. One can ask what fields have this property? If our theory has full quantifier elimination, then it is necessarily the case that models are algebraically closed fields, so the more interesting question is what fields have this property and do not eliminate quantifiers. We say a field K has *Serre's Property (F)* if for every $n \in \mathbb{N}$ there are at most finitely many extensions K' of degree n (eqiv. abs Galois group is small i.e. profinite?). Thus, the related question is: "Does a field K have Serre's Property (F) if and only if K (the theory of K?) has "almost" quantifier elimination?

Chapter V

Strongly Minimal Pseudofinite Structures

In this chapter, we study pseudofinite theories which are also strongly minimal. Strong minimality is a property that one usually first encounters in the theory of algebraically closed fields. It turns out that, in general, strongly minimal theories are very nice in a model theoretic sense. They are, for example, uncountably categorical and have many features which make them geometically interesting as well.

Definition V.1. Let T be a complete, 1-sorted theory. T is called *strongly minimal* if, for every model $M \models T$ and every definable (possibly with parameters) set $X \subseteq M$, X is either finite or cofinite.

We will also say that a structure M is strongly minimal if $\mathrm{Th}(M)$ is strongly minimal. Equivalently, for every elementary extension M' of M and every definable $X \subseteq M'$, X is either finite or cofinite. Note that in a many sorted theory, it makes sense to talk about the sorts being strongly minimal. We will sometimes write D for some saturated, strongly minimal set.

Example V.2. Some examples of strongly minimal theories/structures are:

- ACF_p, the theory of algebraically closed fields of charachteristic p for $p = 0$ or prime;
- $\mathrm{Th}(\mathbb{Q}, +)$, the theory of rational numbers with addition;
- $\mathrm{Th}(\mathbb{F}_p^\omega, +)$;
- $\mathrm{Th}(\mathbb{Z}, s)$, the theory of the integers with a successor function;
- $\mathrm{Th}(\mathbb{N}, s)$, the theory of the natural numbers with a successor function.
- Let G be a group and let X be an infinite, free G-set. Let our language be $L = \{f_g : g \in G\}$. Then $(X, f_g)_{g \in G}$ is a strongly minimal structure, where we interpret the symbols f_g as functions $f_g(x) = gx$.

We will assume, for now, that T is 1-sorted and strongly minimal.

Exercise V.3. For any formula $\varphi(x, \bar{y})$, there is $N_\varphi \in \mathbb{N}$ such that for every $M \models T$ and every $\bar{b} \in M$, we have that $\varphi(x, \bar{b})(M)$ is infinite if it has cardinality greater than or equal to N_φ. This is related to the elimination of the quantifier $(\exists^\infty)$. A theory T eliminates the $(\exists^\infty)$ quantifier if for every formula $\varphi(\bar{x}, \bar{y})$, there is $N_\varphi \in \mathbb{N}$ such that for every $M \models T$ and ever $\bar{b} \in M$, if $|\varphi(\bar{x}, \bar{b})(M)| \geq N_\varphi$ then $|\varphi(\bar{x}, \bar{b})(M)| \geq \omega$.

Lemma V.4. *Let $M \models T$, $A \subseteq M$ and $b_1, b_2 \in M \backslash \operatorname{acl}(A)$. Then* $\operatorname{tp}_M(b_1/\operatorname{acl}(A)) = \operatorname{tp}_M(b_2/\operatorname{acl}(A))$.

Proof. For each $\varphi(x) \in L_A$, either $\varphi(x)(M)$ is finite or cofinite, since T is strongly minimal. If $\varphi(x)(M)$ is finite then $\varphi(x)(M) \subseteq \operatorname{acl}(A)$, so $M \models \neg\varphi(b)$ for every $b \in M \setminus \operatorname{acl}(A)$. That is, for every $b \in M \setminus \operatorname{acl}(A)$ and every $\varphi(x) \in L_A$, $M \models \varphi(b)$ if and only if $\varphi(x)(M)$ is cofinite. Note that if M is saturated enough compared to A, there is some $b \in M \setminus \operatorname{acl}(A)$. $\square$

We remark that, for general T not necessarily strongly minimal, if $M \preccurlyeq \mathcal{N}$, then $\operatorname{acl}_{\mathcal{N}}(M) = M$. Suppose $b \in \mathcal{N}$ and $\mathcal{N} \models \varphi(b, \bar{a})$ where $\bar{a} \in M$ and that $\mathcal{N} \models \exists^{=k} x \varphi(x, \bar{a})$ (i.e. b is algebraic over M). Then $M \models \exists^{=k} x \varphi(x, \bar{a})$ and so all of the (finitely many) realizations of $\varphi(x, \bar{a})$ in $\mathcal{N}$ must already be in M.

Back in the strongly minimal setting, acl has some aditional useful properties:

Lemma V.5 (Symmetry of acl)**.** *Let $M \models T$, $A \subseteq M$ and $b, c \in M$. Suppose that $c \in \operatorname{acl}(A \cup \{b\}) \setminus \operatorname{acl}(A)$. Then $b \in \operatorname{acl}(A \cup \{c\}) \setminus \operatorname{acl}(A)$.*

Proof. Our assumption that $c \in \operatorname{acl}(A \cup \{b\})$ is witnessed by a formula $\varphi(x, y) \in L_A$ such that $M \models \psi(b, c)$ where $\psi(x, y)$ is

$$\varphi(x, y) \wedge \exists^{=k} z \varphi(x, z).$$

We claim that the formula $\psi(x, c)$ (which is realised by b) has finitely many solutions, which shows that b is algebraic over $A \cup \{c\}$. If not, then $\psi(x, c)$ has cofinitely many solutions in M. Let us assume that all but $r \in \mathbb{N}$ many elements of M realise $\psi(x, c)$. Consider the formula $\chi(y) \in L_A$ expressing the the statement "all but r-many elements x satisfy $\psi(x, y)$". By construction, we have that $M \models \chi(c)$. Now, we are also assuming that $c \notin \operatorname{acl}(A)$ and so we can choose $k+1$ distinct realizations $c_1, \ldots, c_{k+1} \in M$ of $\chi(y)$. Therefore, for each c_i, there are only r-many elements $d \in M$ such that $M \models \neg\psi(d, c_i)$. Hence, we can find $b' \in M$ such that

$$M \models \bigwedge_{i=1}^{k+1} \psi(b, c_i).$$

This is a contradiction, however, since $\psi(b, x)$ expresses that there are precisely k-many realizations. This shows that $b \in \operatorname{acl}(A \cup \{c\})$. This completes the proof, since if $b \in \operatorname{acl}(A)$, then $\operatorname{acl}(A \cup \{b\}) = \operatorname{acl}(A)$ and therefore $c \in \operatorname{acl}(A)$, which we assumed was not the case. $\square$

Definition V.6. A set $\{b_i\}_{i\in I}$ is said to be algebraically independent over a set A if $b_i \notin \mathrm{acl}(A, b_j, j \in I \setminus \{i\})$.

Lemma V.7. *Let T be strongly minimal and let $M \models T$.*

1. *Let $A \subseteq M$ and let $\{b_i : i < \alpha\}$ for α an ordinal be such that $b_0 \notin \mathrm{acl}(A)$ and for every $i < \alpha$, $b_i \notin \mathrm{acl}(A, \{b_j : j < i\})$. Then $\{b_i : i < \alpha\}$ is algebraically independent over A.*
2. *Let $B \subset M$ and $A \subseteq M$. Let $B_0 \subseteq B$ be a maximal algebraically independent over A subset of B. Then $B \subseteq \mathrm{acl}(B_0 \cup A)$.*
3. *Let $B \subset M$, $A \subseteq M$. The cardinality of any maximal algebraically independent over A subset of B is unique and denoted* $\dim(B/A)$ *or* $\dim(\bar{b}/A)$ *where $\bar{b}$ is some (any) enumeration of B.*

Proof. 1. Suppose not. Then for some $i < \alpha$,

$$b_i \in \mathrm{acl}(A, b_0, \ldots, \hat{b_i}, \ldots).$$

Let $\gamma < \alpha$ be least such that for any $\varphi(x, \bar{c})$ realizing the algebraicity of b_i, γ is the largest such that b_γ appears as a parameter in $\varphi(x, \bar{c})$. Thus,

$$b_i \in \mathrm{acl}(A, b_0, \ldots, \hat{b_i}, \ldots, b_\gamma) \setminus \mathrm{acl}(A, b_0, \ldots, \hat{b_i}, \ldots, b_\beta : \beta < \gamma).$$

By Lemma V.5,

$$b_\gamma \in \mathrm{acl}(A, b_0, \ldots, b_i, \ldots, b_\beta : \beta < \gamma) \setminus \mathrm{acl}(A, b_0, \ldots, \hat{b_i}, \ldots, b_\beta : \beta < \gamma),$$

but then in particular,

$$b_\gamma \in \mathrm{acl}(A, b_0, \ldots, b_i, \ldots, b_\beta : \beta < \gamma),$$

a contradiction. Note that the converse is immediate.

2. Suppose not. Then there is $b \in B$ such that $b \notin \mathrm{acl}(B_0 \cup A)$. By part 1., $B_0 \cup \{b\}$ is algebraically independent over A. This contradicts the maximality of B_0.

3. Let $C = \{c_0, c_1, \ldots\}$ and $D = \{d_0, d_1, \ldots\}$ be two distinct, maximal, algebraically independent over A subsets of $B \subseteq M$ (with no particular enumeration). Suppose C has cardinality $\lambda \geq \aleph_0$ and D has cardinality κ and that $\lambda > \kappa$. By part 2., $D \subseteq B \subseteq \mathrm{acl}(A, C)$ and so for every $d_i \in D$, there is a formula $\varphi_i(x)$ over $D \cup A$ realizing the algebraicity of d_i. Let E_i be the finite subset of $D \cup A$ appearing as parameters in $\varphi_i(x)$. Let

$$E = \bigcup_{i\in\kappa} E_i.$$

Then, by set theory, $|E| = \kappa < \lambda$ and so there is some $j \in \lambda$ such that $c_j \in C$ does not appear in any E_i. Since D is a maximal algebraically independent subset of B over A, $c_j \in \operatorname{acl}(A, D)$. Then, c_j is algebraic over some finite subset of $A \cup D$. Since every element of D is algebraic over some finite subset of $A \cup E$, we have that $c_j \in \operatorname{acl}(A, E)$. Since $c_j \notin E$, and $E \subsetneq C$, we have contradicted the fact that C is algebraically independent.

Suppose now C and D are maximally algebraically independent over A sets such that C has cardinality $n < \aleph_0$ and, for a contradiction, that D has cardinality $m > n$. We know by maximality that

$$D \subseteq \operatorname{acl}(A, C) \setminus \operatorname{acl}(A).$$

Let $C = \{c_1, \ldots, d_n\}$ and $D = \{d_1, \ldots, d_m\}$. By maximality, there is some d_i and c_j such that

$$d_i \in \operatorname{acl}(A, c_1, \ldots, c_n) \setminus \operatorname{acl}(A, c_1, \ldots, c_{j-1}, c_{j+1}, \ldots, c_n),$$

since otherwise we would have $D \subseteq \operatorname{acl}(A, C \setminus \{c_j\})$ and by maximality of D, C would not be independent. By renumbering if necessary, we may assume d_1 and c_1 are such a pair. By V.5,

$$c_1 \in \operatorname{acl}(A, d_1, c_2, \ldots, c_n) \setminus \operatorname{acl}(A).$$

Since D is an algebraically independent set, $d_2 \notin \operatorname{acl}(A, d_1)$. However, since $c_1 \in \operatorname{acl}(A, d_1, c_2, \ldots, c_n)$, and $d_2 \in \operatorname{acl}(A, C) \setminus \operatorname{acl}(A)$, we have that

$$d_2 \in \operatorname{acl}(A, d_1, c_2, \ldots, c_n) \setminus \operatorname{acl}(A, d_1)$$

and so by symmetry again, we may suppose that

$$c_1, c_2 \in \operatorname{acl}(A, d_1, d_2, c_3, \ldots, c_n) \setminus \operatorname{acl}(A).$$

Repeating this process, we find that

$$c_1, \ldots, c_n \in \operatorname{acl}(A, d_1, \ldots, d_n) \setminus \operatorname{acl}(A).$$

However, since C is maximally algebraically independent over A, we have that there is some $d_k \in D \setminus \{d_1, \ldots, d_n\}$ such that $d_k \in \operatorname{acl}(A, C)$ and hence $d_{n+1} \in \operatorname{acl}(A, d_1, \ldots, d_n)$, which contradicts the fact that D is algebraically independent. Thus, $m \leq n$. By the same argument, $m = n$.

$\square$

Note that, in particular, for any $A \subseteq M$ and any n-tuple $\bar{b}$, $0 \leq \dim\left(\bar{b}/A\right) \leq n$ and, in fact, $\dim\left(\bar{b}/A\right) = \operatorname{RM}\left(\operatorname{tp}_M\left(\bar{b}/A\right)\right)$. The following fact will be useful.

Fact V.8. *Let T be a strongly minimal theory and M a monster model of T. Then for every formula $\varphi(\bar{x}, \bar{y})$, and every $k \in \omega$, the set*

$$\{\bar{b} \in M : \operatorname{RM}(\varphi(\bar{x}, \bar{b})) = k\}$$

is definable in M.

Exercise V.9. 1. Let $\bar{b} = b_1, \ldots, b_n$ and $\bar{c} = c_1, \ldots, c_n$ each be algberaically independent n-tuples over A. Then $\operatorname{tp}\left(\bar{b}/A\right) = \operatorname{tp}\left(\bar{c}/A\right)$.

2. Let $\bar{b}$, $\bar{c}$ be finite tuples, and $A \subseteq M$. Then

$$\dim\left(\bar{b}, \bar{c}/A\right) = \dim\left(\bar{b}/A, \bar{c}\right) + \dim\left(\bar{c}/A\right).$$

Definition V.10. 1. Let $\bar{b}$ be a finite tuple and $A, C \subseteq M$. We say that " $\bar{b}$ is indepenedent from C over A", denoted by $\bar{b} \mathop{\smile\!\!\!\!\!\!|}_A C$, if $\dim\left(\bar{b}/A\right) = \dim\left(\bar{b}/A \cup C\right)$.

2. Let $\bar{b}$ be a possibly infinite tuple. We say that $\bar{b}$ is independent from C over A if $\bar{b}' \mathop{\smile\!\!\!\!\!\!|}_A C$ for all finite subtuples $\bar{b}' \subseteq \bar{b}$.

Equivalently, $\bar{b} \mathop{\smile\!\!\!\!\!\!|}_A C$ if any finite subtuple $\bar{b}'$ of $\bar{b}$ which is algebraically independent over A remains algebraically independent over $A \cup C$. We remark that this notion of independence is precisely non-forking independence in a special case. That is, in a strongly minimal theory T, $\bar{b} \mathop{\smile\!\!\!\!\!\!|}_A C$ if and only if $\operatorname{tp}\left(\bar{b}/A \cup C\right)$ does not divide over C.

It is useful to note that for infinite tuple $\bar{b}$, by definition $\bar{b} \mathop{\smile\!\!\!\!\!\!|}_A C$ if and only if $\bar{b}' \mathop{\smile\!\!\!\!\!\!|}_A C$ for all finite subtuples $\bar{b}' \subset \bar{b}$, and also, if and only if $\bar{b}' \mathop{\smile\!\!\!\!\!\!|}_A C'$ for all finite subtuples $\bar{b}' \subset \bar{b}$ and all finite subsets $C' \subseteq C$.

Exercise V.11. 1. (Transitivity.) Let $A \subseteq B \subseteq C$ and let $\bar{e}$ be a (possibly infinite) tuple. Then $\bar{e} \mathop{\smile\!\!\!\!\!\!|}_A C$ if and only if $\bar{e} \mathop{\smile\!\!\!\!\!\!|}_A B$ and $\bar{e} \mathop{\smile\!\!\!\!\!\!|}_B C$.

2. (Symmetry.) Let $\bar{b}, \bar{c}$ be (possibly infinite) tuples. Then $\bar{b} \mathop{\smile\!\!\!\!\!\!|}_A \bar{c}$ if and only if $\bar{c} \mathop{\smile\!\!\!\!\!\!|}_A \bar{b}$. For $\bar{b}, \bar{c}$ finite tuples, this is the same as showing $\dim\left(\bar{b}, \bar{c}/A\right) = \dim\left(\bar{b}/A\right) + \dim\left(\bar{c}/A\right)$.

3. (Existence of Non-Forking Extensions.) Suppose we are given a tuple $\bar{c}$ and $A \subset B$. Then there exists a tuple $\bar{c}'$ such that $\operatorname{tp}\left(\bar{c}'/A\right) = \operatorname{tp}\left(\bar{c}/A\right)$ and $\bar{c}' \mathop{\smile\!\!\!\!\!\!|}_A B$. We say that $\operatorname{tp}\left(\bar{c}'/B\right)$ is a "*non-forking extension*" of $\operatorname{tp}\left(\bar{c}/A\right)$.

Of particular interest are so called *stationary types*:

Definition V.12. A complete type $p(x) \in S(A)$ is said to be stationary if it has a unique non-forking extension over any set $B \supseteq A$.

Remark V.13. Stationary types in ACF are the generic types of absolutely irreducible varieties.

The following result is true more generally in all stable theories, though we will give a proof for T strongly minimal:

Proposition V.14. *Let T be strongly minimal and let M be a (small) model of T. Let $p\left(\bar{x}\right) \in S(M)$. Then $p\left(\bar{x}\right)$ is stationary.*

Proof. Let $C \supseteq M$ be some set and let $\bar{d} \in \mathcal{U}$ in some very saturated elementary extension of M realize $p(\bar{x})$ so that $p(\bar{x}) = \operatorname{tp}(\bar{d}/M)$. Suppose also that $\bar{d} \mathop{\smile\!\!\!\!\!\!|}_M C$. That is, $\operatorname{tp}(\bar{d}/C)$ is a non-forking extensions of $p(\bar{x})$. Let us write $\bar{d} = \bar{a}\bar{b}$ (we will consider the variable $\bar{x} = \bar{z}\bar{y}$) where $\bar{a}$ is algebraically independent over M and $\bar{b} \in \operatorname{acl}(\bar{a}, M)$. It suffices to show that for $\bar{a}'\bar{b}'$ realizing $p(\bar{x})$ such that $\bar{a}'\bar{b}' \mathop{\smile\!\!\!\!\!\!\mid}_M C$, it must be the case that $\operatorname{tp}(\bar{a}\bar{b}/C) = \operatorname{tp}(\bar{a}'\bar{b}'/C)$. Suppose we have such a $\bar{a}'\bar{b}'$. Since $\bar{b} \in \operatorname{acl}(\bar{a}, M)$, the type $\operatorname{tp}(\bar{b}/M\bar{a})$ is algebraic. Suppose that this type has k-many realizations. By compactness, there is a formula $\varphi(\bar{a}, \bar{y})$ over M which isolates $\operatorname{tp}(\bar{b}/M\bar{a})$. Since $\varphi(\bar{z}\bar{y}) = \varphi(\bar{x}) \in p(\bar{x})$, we have that $\models \varphi(\bar{a}', \bar{b}')$. Now, by Exercise V.9, $\bar{a}$ and $\bar{a}'$ have the same type over M, so take some automorphism σ fixing M which sends $\bar{a}$ to $\bar{a}'$. Then $\varphi(\bar{a}', \bar{y})$ isolates the type $\operatorname{tp}(\sigma(\bar{b})/M\bar{a}') = \operatorname{tp}(\bar{b}'/M\bar{a}')$. Now, if $\varphi(\bar{a}, \bar{y})$ isolates a complete type over C, we are done. This is because, $\bar{a}'\bar{b}' \mathop{\smile\!\!\!\!\!\!\mid}_M C$ if and only if $\bar{a}'$ is algerbaically independent over C and so by Exercise V.9.1 $\operatorname{tp}(\bar{a}/C) = \operatorname{tp}(\bar{a}'/C)$. By similar reasoning as before, $\operatorname{tp}(\bar{b}/C\bar{a}) = \operatorname{tp}(\bar{b}'/C\bar{a}')$ and so $\operatorname{tp}(\bar{a}\bar{b}/C) = \operatorname{tp}(\bar{a}'\bar{b}'/C)$.

For a contradiction, let us suppose that $\varphi(\bar{a}, \bar{y})$ does not isolate a complete type over C. Then, it must be the case that there is $\psi(\bar{a}, \bar{c}, \bar{y})$ over M, with $\bar{c} \in C$ such that

- $\models \psi(\bar{a}, \bar{c}, \bar{b})$;
- $\models \forall \bar{y}(\psi(\bar{a}, \bar{c}, \bar{y}) \to \varphi(\bar{a}, \bar{y}))$;
- $\psi(\bar{a}, \bar{c}, \bar{y})$ has $k' < k$ many solutions.

That is, $\psi(\bar{a}, \bar{c}, \bar{y})$ isolates a complete algebraic type and the type $\operatorname{tp}(\bar{a}/C)$ is definable over M. Let $\chi(\bar{z})$ be the definition over M (in fact, over $\emptyset$) of the formula

$$\forall \bar{y}\left(\psi\left(\bar{x}, \bar{z}, \bar{y}\right) \to \varphi\left(\bar{x}, \bar{y}\right)\right) \wedge \exists^{=k'} \bar{y} \psi\left(\bar{x}, \bar{z}, \bar{y}\right).$$

Now, $\models \chi(\bar{c})$ and so there is $\bar{c}' \in M$ such that $M \models \chi(\bar{c}')$. This contradicts the fact that $\varphi(\bar{a}, \bar{y})$ isolates $\operatorname{tp}(\bar{b}/M\bar{a})$.

□

Corollary V.15. *Let T be strongly minimal and let $M \models T$. Let $\bar{b}_1, \bar{b}_2$ in some saturated elementary extension of M be such that $\bar{b}_1 \mathop{\smile\!\!\!\!\!\!\mid}_M \bar{b}_2$. Then* $\operatorname{tp}(\bar{b}_1, \bar{b}_2/M)$ *is determined by* $\operatorname{tp}(\bar{b}_1/M)$ *and* $\operatorname{tp}(\bar{b}_2/M)$.

Proof. Exercise. □

Definition V.16. Let D be a saturated, strongly minimal set and let $\varphi(\bar{x})$ be a formula defining $X \subseteq D^n$ with parameters from some small set $A \subset D$. Then:

1. $\dim(\varphi(\bar{x})) = \dim(X) = \max\{\dim(\bar{b}/A) : \bar{b} \in X \subseteq D^n\}$;
2. if $\dim(X) = k$, define the *degree* or *multipicity* of X, $\operatorname{mult}(X)$ as the maximum d such that one can write $X = X_1 \sqcup \ldots \sqcup X_d$, a disjoint union of definable sets X_i, such that $\dim(X_i) = k$ for $1 \leq i \leq d$.

Exercise V.17. The definition of dimension does not depend on A. The dimension is the same if one replaces A with small $B \supset A$.

Note that this definition of dimension and degree aligns well with the corresponding definitions for varieties.

Lemma V.18. *Suppose $\bar{b} \in D^n$ is such that $\dim(\bar{b}/A) = k$. Then there exists a formula $\psi(\bar{x}) \in \mathrm{tp}(\bar{b}/A)$ such that $\dim(\varphi(\bar{x})) = k$.*

Proof. We may assume that $\bar{b} = (b_1, \ldots, b_k, b_{k+1}, \ldots, b_n)$ is such that $(b_1, \ldots, b_k)$ are algebraically independent over A. Hence, $b_{k+1}, \ldots, b_n \in \mathrm{acl}(A, b_1, \ldots, b_k)$. This is witnessed by some formula $\varphi(x_1, \ldots, x_k, x_{k+1}, \ldots, x_n)$ over A such that $D \models \varphi(\bar{b})$ and

$$D \models \forall (x_1, \ldots, x_k) \exists^{\leq r} (x_{k+1}, \ldots, x_n) \varphi(\bar{x})$$

for some $r \in \mathbb{N}$. Therefore, $\dim(\varphi(\bar{x})) \leq k$. Since $\dim(\bar{b}/A) = k$ and $\varphi(\bar{b})$, we have $\dim(\varphi(\bar{x})) = k$. $\square$

Lemma V.19. *1.* $\dim(D^n) = n$;

2. $\mathrm{mult}(D^n) = 1$.

Proof. 1. Given any $n \in \mathbb{N}$, by saturation of D, we can find $a_i \in D$ such that $a_i \notin \mathrm{acl}(a_1, \ldots, a_{i-1})$ for $i \leq n$. Therefore $(a_1, \ldots, a_n) \in D^n$ is an algerbaically independent tuple of maximum lenght n. Hence $\dim(D^n) = n$.

2. Suppose, for a contradiction, that $D^n = X \sqcup Y$ where X and Y are definable (over some set A) and that $\dim(X) = \dim(Y) = n$. Then we may choose $\bar{b} \in X$ and $\bar{c} \in Y$ with $\dim(\bar{b}/A) = \dim(\bar{c}/A) = n$. However, $\mathrm{tp}(\bar{b}/A) \neq \mathrm{tp}(\bar{c}/A)$ since the formula $x \in X$ is in $\mathrm{tp}(\bar{b}/A)$ and not in $\mathrm{tp}(\bar{c}/A)$. This contradicts the uniqueness of the non-algebraic type (Exercise V.9).

$\square$

Definition V.20. By a *k-cell*, we mean a definable set $X \subseteq D^n$ for some $n \geq k$ such that under some projection $\pi : D^n \to D^k$, we have $\dim(\pi(X)) = k$ and $\pi \restriction_X : X \twoheadrightarrow \pi(X)$ is r-to-1 for some $r > 0$.

Note that any k-cell X definable over a set A has dimension k and, for every $\bar{b} \in X$, $\bar{b} \in \mathrm{acl}(\pi(\bar{b}), A)$.

Proposition V.21. *Let $X \subseteq D^n$ be definable. Then X is a finite union of cells.*

Proof. Suppose that X is defined over A. Let $\bar{b} \in X$ and suppose $\dim(\bar{b}/A) = k$ with $(b_1, \ldots, b_k)$ algebraically independent over A. Let $r = \mathrm{mult}(\mathrm{tp}(\bar{b}/Ab_1, \ldots, b_k))$ (i.e. the number of realizations of $\mathrm{tp}(\bar{b}/Ab_1, \ldots, b_k)$, since each realization is a point of dimension k and there are only finitely many since the type is algebraic). Let $\varphi_{\bar{b}}(\bar{x}) \in \mathrm{tp}(\bar{b}/Ab_1, \ldots, b_k)$ be a formula witnessing this. That is,

$$D \models \forall x_1, \ldots, x_k \left(\exists x_{k+1}, \ldots, x_n \varphi_{\bar{b}}(\bar{x}) \leftrightarrow \exists^{=r} x_{k+1}, \ldots, x_n \varphi_{\bar{b}}(\bar{x}) \right)$$

and

$$\dim\left(\exists x_{k+1},\ldots,x_n\varphi_{\bar{b}}(\bar{x})\right)=k.$$

It is immediate from the definition that $\varphi_{\bar{b}}(\bar{x})$ defines a k-cell, since the projection of $\varphi_{\bar{b}}(\bar{x})(D)$ onto the cooridnates $x_{k+1},\ldots,x_n$ has dimension k and the fibres are of size r. Therefore,

$$\text{“}x\in X\text{”}\models\bigvee_{\bar{b}\in X}\varphi_{\bar{b}}(\bar{x}).$$

The result follows by compactness by appealing to the fact that the $\varphi_{\bar{b}}(\bar{x})$ depends only on the type of $\bar{b}$ over A and so there are few enough types since A is small. □

Remark V.22. 1. In the proof of Proposition V.21, we may assume that $A=M\preccurlyeq D$ is a model and argue as in the proof of Lemma V.14 to show that each of the cells is of multiplicity 1. That is, we let $\bar{b}\in X$ defined over M and consider $\mathrm{tp}(\bar{b}/M)$ with $b_1,\ldots,b_k$ algebraically independent over M and $\bar{b}\in\mathrm{acl}(b_1,\ldots,b_k)$. Choose $\psi_{\bar{b}}(\bar{x})$ such that $\psi_{\bar{b}}(b_1,\ldots,b_k,x_{k+1},\ldots,x_n)$ isolates $\mathrm{tp}(b_{k+1},\ldots,b_n/Mb_1,\ldots b_k)$. Then $\psi_{\bar{b}}(\bar{x})$ has dimension k and multiplicity 1, as in the proof of V.14. Thus, $\mathrm{mult}(X)$ exists for any definable X.

2. We can also require that X is a finite disjoint union of cells.

Example V.23. Let $D=\mathbb{C}$ and let $X\subset D^2$ be $X=\{(x,y):x=y^2\wedge x\neq 0\}$. Let $\pi:X\to\mathbb{C}$ be the projection onto the x-coordinate. Then $\pi(X)=\mathbb{C}\setminus\{0\}$, which has dimension 1. Note that $\pi\restriction_X$ is a 2-to-1 map. X has dimension 1 and multiplicity 1, since it is an irreducible curve (X is a 1-cell).

Lemma V.24. *Let X,Y be definable over a set A with dimension k_1 and k_2 respectively and both of multiplicity 1. Then* $\dim(X\times Y)=k_1+k_1$ *and* $\mathrm{mult}(X\times Y)=1$.

Proof. Since X and Y are definable over A, so is $X\times Y$. Let $(\bar{b}_1,\bar{b}_2)\in X\times Y$. We have $\dim(\bar{b}_i/A)\leq k_i$ for $i=1,2$ and so $\dim(\bar{b}_1,\bar{b}_2/A)\leq k_1+k_2$ and hence $\dim(X\times Y)\leq k_1+k_2$. For equality, choose $\bar{b}_1\in X$ and $\bar{b}_2\in Y$ such that $\dim(\bar{b}_i/A)=k_i$ for $i=1,2$ and furthermore, such that $\bar{b}_1\downarrow_A\bar{b}_2$. Then $\dim(\bar{b}_1,\bar{b}_2/A)=\dim(\bar{b}_1/A)+\dim(\bar{b}_2/A)=k_1+k_2$.

It remains to show that $\mathrm{mult}(X\times Y)=1$. Observe that if $\mathrm{mult}(X\times Y)>1$, then there would be $Z_1,Z_2\subset X\times Y$ definable over some model $M\supset A$ such that $Z_1\cap Z_2=\emptyset$ and $\dim(Z_i/M)=k_1+k_2$ for $i=1,2$. Therefore, we can find $\bar{c}_1\in Z_1$ and $\bar{c}_2\in Z_2$ such that $\dim(\bar{c}_i/M)=k_1+k_2$ for $i=1,2$ but $\mathrm{tp}(\bar{c}_1/M)\neq\mathrm{tp}(\bar{c}_2/M)$. Therefore, it suffices to show that for any $M\supseteq A$, M a model, there exists a unique type over M of dimension k_1+k_2 containing the formula “$(x,y)\in X\times Y$”. So, let M be such a model, $(\bar{b}_1,\bar{b}_2)\in X\times Y$ with $\dim(\bar{b}_1,\bar{b}_2/M)=k_1+k_2$. This forces:

1. $\dim(\bar{b}_1/M) = k_1$, since otherwise $\dim(\bar{b}_1/M) < k_1$ and so $\dim(\bar{b}_2/M) > k_2$, which is impossible;
2. $\dim(\bar{b}_2/M) = k_2$, for the same reason as above;
3. $\bar{b}_1 \mathop{\smile\!\!\!\!\!\!\!\;\mid}_M \bar{b}_2$, since if not, $\dim(\bar{b}_1, \bar{b}_2/M) \neq \dim(\bar{b}_1/M) + \dim(\bar{b}_2/M) = k_1 + k_2$.

By (i) and (ii), as X and Y each have multiplicity 1, $\mathrm{tp}(\bar{b}_1/M)$ and $\mathrm{tp}(\bar{b}_2/M)$ are uniquely determined (by the same observation that we made assuming $\mathrm{mult}(X \times Y) > 1$). BY (iii) and Corollary V.15, $\mathrm{tp}(\bar{b}_1, \bar{b}_2/M)$ is uniquely determined. $\square$

One might ask what can be said about a theory T which is both strongly minimal and pseudofinite. Traditionally, theories which were categorical (more specifically, $\aleph_0$-categorical) were more popular as a topic of study that pseudofinite theories. For a strongly minimal theories, we have the following automatically:

Proposition V.25. *If T is strongly minimal, then T is uncountably categorical.*

Proof. $\square$

However, $\aleph_0$-categoricity is not automatic and was a natural hypothesis. The following big theorem was proven by Zilber [13] and independently by Cherlin, Harrington, and Lachlan [4]:

Theorem V.26. *Any $\aleph_0$-categorical, strongly minimal (complete) theory T in a countable language is pseudofinite.*

In fact, the original conclusion was stated as "T is not finitely axiomatizable," as the property of being pseudofinite was not seen as being particularly important or interesting, but the proof goes by pseudofiniteness.

Note that, originally, Zilber, Cherlin, Harrington, and Lachlan proved Theorem V.26 in order to show that theories satisfying the hypotheses were not finitely axiomatizable. That result is a direct consequence of V.26, since any finitely axiomatizable, pseudofinite theory has a finite model. It in particular is axiomatized by a single sentence (the conjunction of the axioms), and that sentence must have a finite model by pseudofiniteness. This fact is mainly of historical interest, since pseudofiniteness is now considered more important than finite axiomatizability.

Geometric Notions, Again

We continue our study of the geometric notions related to strong minimality and pseudofiniteness.

Definition V.27. A strongly minimal, complete theory T is locally modular if, in a saturated model D, we have, for all finite tuples a and b from D^{eq},

$$\dim(ab) = \dim(a) + \dim(b) - \dim(\mathrm{acl}(a) \cap \mathrm{acl}(b)).$$

Local modularity generalizes the corresponding property of dimensions in linear algebra. The notion yields a nice structural result:

Remark V.28. Let T be a locally modular (also strongly minimal and complete). Then exactly one of the following holds of D, a saturated model:

1. acl is trivial: for $A \subseteq D$,
$$\operatorname{acl}(A) = \bigcup_{a \in A} \operatorname{acl}(a)$$

2. There is a strongly minimal interpretable (i.e., definable in T^{eq}) commutative group G in D. Moreover, every definable subset of G^n (for all $n \in \omega$) is a finite boolean combination of translates of definable subgroups.

Note that case 1 essentially says that all the interesting relations are binary (for an appropriate sense of "interesting"). Examples of case 1 include a set with no structure and the theory the integers with the successor function.

Exercise V.29. ACF_p is not locally modular for any p.

Remark V.28 is largely the content of Hrushovski's PhD thesis [5]. Boris Zilber showed the same conclusion for $\aleph_0$-categorical strongly minimal theories. Also note that in [13],[4], they probe that strongly minimal $\aleph_0$-categorical theories are locally modular. Cherlin, Harrington, and Lachlan's proof uses the classification of finite simple groups, and Zilber used model theoretic methods. See [9] for details.

It is an open problem to find a similar structure result for strongly minimal theories in general.

In the late 1980s, Hrushovski [6] gave an account of Zilber's proof in more generality by proving that, for a complete strongly minimal theory,

- $\aleph_0$-categoricity implies an additional property called "unimodularity" (fairly easy), and
- unimodularity implies local modularity.

Definition V.30. Let D be a saturated and strongly minimal. D is unimodular if, whenever a and b are interalgebraic, algebraically independent tuples of the same finite length n, the multiplicity of a over b equals the multiplicity of b over a.

Observe that any $\aleph_0$-categorical strongly minimal theory is unimodular.

Exercise V.31. Prove that $\aleph_0$-categorical, strongly minimal theories are unimodular. (Hint: use the fact that, in an $\aleph_0$-categorical theory, the algebraic closure of a finite set is finite, and use a counting argument.)

We will eventually prove in these notes that strongly minimal, pseudofinite structures are unimodular (and hence locally modular). By V.26, this generalizes the work of Zilber / Cherlin / Harrington / Lachlan.

For the rest of this chapter, let D be a strongly minimal, pseudofinite, saturated model. We will consider D equiped with the general non-standard counting machinery from, e.g., Construction I.9.

Proposition V.32. *Suppose $|D| = q \in \mathbb{N}^*$. Let $X \subseteq D^n$ be definable (with parmaters). Then there is a polynomial $P_X \in \mathbb{Z}[x]$ with positive leading coefficient such that $|X| = P_x(q)$. Moreover, the degree of P_X is the Morley rank of X.*

Moreover, given $\varphi(\bar{x}, \bar{y}) \in L$, there exists $\psi_1(\bar{y}) \dots \psi_k(\bar{y}) \in L$ and $P_1 \dots P_k \in \mathbb{Z}[x]$ such that, for any $\bar{b}$, exactly one of $\psi_i(\bar{b})$ holds, and $|\varphi(D, \bar{b})| = P_i(q)$ whenever $\models \psi_i(\bar{b})$.

The key point of Proposition V.32 is that we know the exact cardinality of definable sets, and it's controlled internally (without resorting to the language expansion in, e.g., Construction I.9).

Proof. We leave the proof of the final part of the claim (i.e., definability of the parameters giving a certain size) as an exercise, and proof the rest here.

Note that D has Morley rank 1, and $|D| = q$. Therefore we can take $P_D(x) = x$. More generally, $P_{D^n}(x) = x^n$.

We proceed by induction on $\dim(X)$. If $\dim(X) = 0$, X is finite, so $P_X(x) = m$, where $m = |X| \in \mathbb{Z}$.

For the inductive step, let $X \subseteq D^n$. By Proposition V.21, X is a finite disjoint union of cells. We may therefore assume that $\dim(X) = k$, and X is a k-cell. That is, there is a projection $\pi : D^n \to D^k$ whose image has dimension k, with $\pi \restriction X$ an r-to-1 map.

By Lemma V.19, D^k has dimension k and multiplicity 1, so $\dim(D^k \backslash \pi(X)) < k$. By the inductive hypothesis, $|D^k \backslash \pi(X)| = Q(q)$, where $Q \in \mathbb{Z}[x]$ has positive leading coefficient and degree $< k$. Therefore

$$\begin{aligned} |X| &= r|\pi(X)| \\ &= r\left(|D^k| - |D^k \setminus \pi(X)|\right) \\ &= r\left(q^k - Q(q)\right). \end{aligned}$$

□

Note that, in the case of pseudofinite fields (which are not strongly minimal), the conclusions of Proposition V.32 (in particular getting an exact count) are closely related to the Weil conjectures. Recall that we do have a similar result for fields in Fact IV.8 but that only gives an asymptotic estimate.

The counting result in Proposition V.32 gives us unimodularity. They key to this result is that we can apply similar reasoning to the $\aleph_0$-categorical case: in that case we could count points because the relevant sets were finite; here we can count as a result of Proposition V.32.

Corollary V.33. *D (which is strongly minimal, pseudofinite) is unimodular, and therefore locally modular.*

Proof. Let $\bar{a}, \bar{b}$ be generic, interalgebraic, and length n. Take $\varphi(\bar{x}, \bar{y})$ such that $\varphi(\bar{x}, \bar{b})$ isolates $\mathrm{tp}(\bar{a}/\bar{b})$ and $\varphi(\bar{a}, \bar{x}$ isolates $\mathrm{tp}(\bar{b}/\bar{a})$. Let $k = \mathrm{mult}(\bar{a}/\bar{b})$ and $l = \mathrm{mult}(\bar{b}/\bar{a})$. We want to show that $k = l$.

Let $Z \subseteq D^n \times D^n$ be the (definable without parameters) set of pairs $(\bar{a}_0, \bar{b}_0)$ realizing $\varphi(\bar{x}, \bar{y})$ such that $\exists^{=k}\bar{x}\varphi(\bar{x}, \bar{b}_0)$ and $\exists^{=l}\bar{y}\varphi(\bar{a}_0, \bar{y})$. Note that:

- $(\bar{a}, \bar{b}) \in Z$
- Fibers of left projection of Z are never larger than l and such fibers of algebraically independent tuples are of size exactly l (since all such tuples have the same type as $\bar{a}_0$.
- Similarly for fibers of right projection with respect to size k.

It is clear that $\dim(Z) = n$, so take $s(t)$ a polynomial with rational coefficients such that $|Z| = s(|D|)$; s has degree n. Let $Z_0 \subseteq Z$ be the points for which left-projection is exactly l-to-1, and $X_0 \subseteq D^n$ be their left projections. Note that Z_0 and X_0 are definable without parameters. Since X_0 contains $\bar{a}$, which is algebraically independent, $\dim(X_0) = n$. Furthermore, since the *degree* of D^n is 1, $\dim(D^n \setminus X_0) < n$. Compute:

$$\begin{aligned}
|s(|D|) - l|D|^n| &= ||Z| - l|D|^n| \\
&\leq ||Z_0| - l|D|^n| + |Z \setminus Z_0| \\
&= |l\,|X_0| - l|D|^n| + |Z \setminus Z_0| \\
&\leq l\,||X_0| - |D|^n| + l\,|D^n \setminus X_0| \\
&= 2l\,|D^n \setminus X_0| \\
&< |D|^{n-\frac{1}{2}}.
\end{aligned}$$

Therefore the leading coefficient of s is l. By the same argument, the leading coefficient of s is k. Therefore $k = l$. □

This concludes our discussion of strongly minimal, pseudofinite theories.

Chapter VI

Definable Regularity Lemmas

Recall that the Szemeredi regularity lemma is about subsets of *any* finite graphs, which we know is equivalent to a statement about internal subsets of any pseudofinite graph. In this chapter, we present some stronger theorems about cases where we know that the graph is tame in some model-theoretic sense. For example, we may consider

- graphs definable in a strongly minimal, pseudofinite structure ([10]),
- graphs definable in a pseudofinite field ([12]),
- pseudofinite graphs such that the edge relation is stable ([7])

In these various more specializes cases, we can get stronger results than those in Chapter II. These strengthenings come in two varieties. Some cases give results which eliminate the "exceptional sets" of classes for which regularity fails. Other cases strengthen the regularity to a dichotomy of edges "almost everywhere" or "almost nowhere," for suitable such notions.

Stable Case

We turn first to the case of pseudofinite graphs with a stable edge relation, following [7]. Recall the definition of a stable relation:

Definition VI.1. Let T be a complete theory, $\varphi(\bar{x}, \bar{y})$ a formula, and $k \in \omega$. Then φ is k-stable if there are no sequences

$$\bar{a}_1, \ldots, \bar{a}_k \bar{b}_1, \ldots, \bar{b}_k$$

in any model such that $\models \varphi(\bar{a}_i, \bar{b}_j)$ iff $i < j$. A formula is stable if it is k-stable for some k.

On the face of it, the main result in [7] is about all finite graphs with a k-stable edge relation, for fixed k. We present here an equivalent form of the theorem for pseudofinite graphs with a stable edge relation. The equivalence is via the same mechanism as in Chapter II.

Proposition VI.2. *Let (V, W, E) be a saturated, pseudofinite, bipartite graph whose edge relation E is stable. Fix $\epsilon > 0 \in \mathbb{R}$. Then we can write*

$$V = V_1 \sqcup \ldots \sqcup V_r$$
$$W = W_1 \sqcup \ldots \sqcup W_s,$$

where these sets are internal, such that

$$||V_i| - |V_j|| \leq 1$$
$$||W_i| - |W_j|| \leq 1.$$

We can take this partition to have the property that, for all $i \leq r$ and all $j \leq s$, either

1. *For at least $\epsilon |V_i|$ many $a \in V_i$ there are at least $\epsilon |W_j|$ many $b \in W_j$ with $\models E(a, b)$, or*

2. *For at least $\epsilon |V_i|$ many $a \in V_i$ there are at least $\epsilon |W_j|$ many $b \in W_j$ with $\models \neg E(a, b)$.*

The conclusion of Proposition VI.2 is quite strong. There is no exceptional set of partitions, and we furthermore have that the pairs are not only regular, they are either dense or codense. Note that the partition is still only into internal sets, not definable sets.

The proof of Proposition VI.2 uses the notion of $\varphi-2$ rank in a finite context, among other techniques. It may be possible to find a conceptual proof in the pseudofinite context using local stability, forking, and counting measures, but such efforts so far have not been successful.

Field Case

Here we present a version of Szemeredi regularity for pseudofinite fields. This result is due to Tao [12], with a refinement due to Pillay and Starchenko [11].

Proposition VI.3. *Let L be the language of fields, F a saturated, pseudofinite field, and (V, W, E) a bipartite graph definable in F. The we can write*

$$V = V_1 \sqcup \ldots \sqcup V_r$$
$$W = W_1 \sqcup \ldots \sqcup W_s,$$

where these sets are definable, such that, for some fixed $C \in \mathbb{R}^+$

$$|V_i| \geq \frac{|V|}{C}$$

$$|W_i| \geq \frac{|W|}{C}.$$

We can take this partition to have the property that, for all $i \leq r$ and $j \leq s$, there is d_{ij} such that, for all $A \subseteq V_i$ and $B \subseteq W_j$ internal subsets,

$$||E \cap A \times B| - d_{ij} |A| |B|| \leq C |F|^{-\frac{1}{4}} |V_i| |W_j| ,$$

where C is the same as above.

Note that, as in the stable case, there are no exceptional sets. In contrast to the stable case, we do not get that the sets are either dense or codense, but we do get that the partitions are definable. The regularity we get in the style of weak Szemeredi regularity (Proposition II.4), further strengthened so that the error factor is the infinitessimal $C|F|^{-\frac{1}{4}}$, rather than $\epsilon \in \mathbb{R}^+$.

The proof of Proposition VI.3, which we do not reproduce here in full, operates crucially via the following lemma. To parse that lemma, recall from Theorem IV.25 that in a pseudofinite saturated field F, a set X defined by $\varphi(\bar{x}, \bar{a})$ has a dimension d and measure μ such that

$$\left| |X| - \mu |F|^d \right| < C |F|^{d-\frac{1}{2}}$$

for some $C \in \mathbb{R}^+$. Furthermore, both d and μ depend definably on $\bar{a}$.

Lemma VI.4. *In the context Proposition VI.3, assume* $\dim(V) = n$, $\dim(W) = k$*, and both are defined over some small $A \subseteq F$. Then we can partition W into sets*

$$W_1, \ldots, W_m$$

definable over $\mathrm{acl}(A)$*, all of dimension k, such that, for all i, j, there are $c_{ij} \in \mathbb{Q}^+$ and $D_{ij} \subset W_i \times W_j$ definable over* $\mathrm{acl}(A)$ *of dimension $< 2k$ such that either*

1. *for all $(a, b) \in W_i \times W_j \setminus D_{ij}$,* $\dim(E(x, a) \cap E(x, b)) < n$*, or*
2. *for all $(a, b) \in W_i \times W_j \setminus D_{ij}$,* $(\dim, \mu)(E(x, a) \cap E(x, b)) = (n, c_{ij})$.

Bibliography

[1] Z. Chatzidakis. Notes on the model theory of finite and pseudo-finite fields. `http://www.logique.jussieu.fr/~zoe/papiers/Helsinki.pdf`.

[2] Z. Chatzidakis. Model theory of finite fields and pseudo-finite fields. *Annals of Pure and Applied Logic*, 88(23):95 – 108, 1997. Joint AILA-KGS Model Theory Meeting.

[3] Z. Chatzidakis, L. van den Dries, and A. Macintyre. Definable sets over finite fields. *Journal fr die reine und angewandte Mathematik (Crelles Journal)*, 427:107–136, 1992.

[4] G. Cherlin, L. Harrington, and A. H. Lachlan. $\aleph_0$-Categorical, $\aleph_0$-Stable Structures. *Annals of Pure and Applied Logic*, 28(2):103–135, 1985.

[5] E. Hrushovski. *Contributions to Stable Model Theory.* PhD thesis, University of California at Berkeley.

[6] E. Hrushovski. Unimodular minimal structures. *J. London Math. Soc*, 2:46, 1992.

[7] M. Malliaris and S. Shelah. $\aleph_0$-Categorical, $\aleph_0$-Stable Structures. *Trans. Amer. Math. Soc.*, 366(3):1551–1585, 2014.

[8] S. Pierre. A Guide to NIP Theories. `http://arxiv.org/abs/1208.3944`, 2014.

[9] A. Pillay. *Geometric Stability Theory.* Oxford logic guides. Clarendon Press, 1996.

[10] A. Pillay. Strongly minimal pseudofinite structures. `http://arxiv.org/abs/1411.5008`, 2014.

[11] A. Pillay and S. Starchenko. Remarks on Tao's algebraic regularity lemma. `http://arxiv.org/abs/1310.7538`, 2013.

[12] T. Tao. Expanding polynomials over finite fields of large characteristic, and a regularity lemma for definable sets. `http://arxiv.org/abs/1211.2894`, 2013.

[13] B.I. Zilber. Totally categorical theories: Structural properties and the non-finite axiomatizability. In Leszek Pacholski, Jedrzej Wierzejewski, and AlecJ. Wilkie, editors, *Model Theory of Algebra and Arithmetic*, volume 834 of *Lecture Notes in Mathematics*, pages 381–410. Springer Berlin Heidelberg, 1980.